LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

The Songs Of Our Past

Visualizing Music Listening Histories

Dissertation

an der Fakultät für Mathematik, Informatik und Statistik
der Ludwig-Maximilians-Universität München

vorgelegt von
Diplom-Medieninformatiker
Dominikus Baur

München, den 5. Oktober 2011

Erstgutachter: Prof. Dr. Andreas Butz
Zweitgutachter: Prof. Dr. Sheelagh Carpendale

Tag der mündlichen Prüfung: 15. Dezember 2011

ABSTRACT

Advancements in technology have resulted in unique changes in the way people interact with music today: Small, portable devices allow listening to it everywhere and provide access to thousands or, via streaming, even millions of songs. In addition, all played tracks can be logged with an accuracy down to the second. So far, these music listening histories are mostly used for music recommendation and hidden from their actual creators. But people may also benefit from this data more directly: as memory extensions that allow retrieving the name of a title, for rediscovering old favorites and reflecting about their lives. Additionally, listening histories can be representations of the implicit relationships between musical items. In this thesis, I discuss the contents of these listening histories and present software tools that give their owners the chance to work with them.

As a first approach to understanding the patterns contained in listening histories I give an overview of the relevant literature from musicology, human-computer-interaction and music information retrieval. This literature review identifies the context as a main influence for listening: from the musical and temporal to the demographical and social.

I then discuss music listening histories as digital memory extensions and a part of lifelogging data. Based on this notion, I present what an ideal listening history would look like and how close the real-world implementations come. I also derive a design space, centered around *time*, *items* and *listeners*, for this specific type of data and shortcomings of the real-world data regarding the previously identified contextual factors.

The main part of this dissertation describes the design, implementation and evaluation of visualizations for listening histories. The first set of visualizations presents listening histories in the context of lifelogging, to allow analysing one's behavior and reminiscing. These casual information visualizations vary in complexity and purpose. The second set is more concerned with the musical context and the idea that listening histories also represent relationships between musical items. I present approaches for improving music recommendation through interaction and integrating listening histories in regular media players.

The main contributions of this thesis to HCI and information visualization are: First, a deeper understanding of relevant aspects and important patterns that make a person's listening special and unique. Second, visualization prototypes and a design space of listening history visualizations that show approaches how to work with temporal personal data in a lifelogging context. Third, ways to improve recommender systems and existing software through the notion of seeing relationships between musical items in listening histories. Finally, as a meta-contribution, the casual approach of all visualizations also helps in providing non-experts with access to their own data, a future challenge for researchers and practitioners alike.

ZUSAMMENFASSUNG

Der technologische Fortschritt der vergangenen Jahre hat das Verhältnis zwischen Menschen und Musik grundlegend geändert: Kleine, tragbare Geräte ermöglichen das Musikhören in jeder Situation und bieten Zugriff auf Tausende oder per Internetanbindung sogar Millionen von Liedern. Zusätzlich kann der Verlauf der abgespielten Lieder problemlos und sekundengenau mitgeloggt werden. Bisher werden dadurch entstandene *Hörhistorien* vor allem für Empfehlungssysteme genutzt und ihren Erzeugern meist vorenthalten, obwohl diese sehr direkt von ihnen profitieren könnten: Als Gedächtniserweiterungen zum schnellen Zugriff auf einzelne Titel, um alte Favoriten wiederzuentdecken oder um über das eigene Leben zu reflektieren. Hörhistorien können auch als Repräsentationen der impliziten Zusammenhänge zwischen einzelnen Liedern genutzt werden. In dieser Dissertation untersuche ich den Inhalt dieser speziellen Datensätze und stelle digitale Werkzeuge vor, die ihren Erzeugern den Zugriff darauf ermöglichen.

Als ersten Ansatz um die Muster, die in derlei Hörhistorien enthalten sind zu verstehen, präsentiere ich einen Überblick über relevante Arbeiten aus den Musikwissenschaften, der Mensch-Maschine-Interaktion und dem Music Information Retrieval. Dieser Literaturüberblick führt zur Feststellung, dass der Hörkontext (vom musikalischen und zeitlichen bis demographischen und sozialen) den Haupteinfluss auf das Hörverhalten darstellt.

Im Anschluß bespreche ich Hörhistorien als digitale Gedächtsniserweiterungen und als Teil des sogenannten Lifeloggings. Basierend auf diesem Konzept stelle ich vor wie eine ideale Hörhistorie aussehen würde und wie nah die realweltlichen Implementierungen diesem Ideal kommen. Ich leite auch einen Designraum basierend auf den Eckpfeilern *Zeit*, *Lieder* und *Hörer* für diesen speziellen Datentyp ab und erkläre, mit welchen Unzulänglichkeiten man bei realen Daten rechnen muss im Zusammenhang mit den vorher identifizierten kontextuellen Faktoren.

Der Hauptteil dieser Dissertation beschreibt das Design, die Implementierung und Evaluation von Visualisierungssystemen für Hörhistorien in zwei generellen Ansätzen: Die ersten Beispiele beschäftigen sich mit Hörhistorien als Teilmenge von Lifelogging-Daten und damit der Analyse des eigenen Verhaltens sowie der Erinnerung an vergangene Erlebnisse. Diese Systeme unterscheiden sich in Komplexität und Zweck und richten sich vorrangig an Laien, die mit ihren persönlichen Daten arbeiten. Die zweite Gruppe von Visualisierungsbeispielen legt ihren Schwerpunkt mehr auf den musikalischen Kontext und basiert darauf, dass Hörhistorien auch für Zusammenhänge zwischen einzelnen Liedern stehen. Dazu präsentiere ich Ansätze um Empfehlungssysteme durch Interaktion zu verbessern und Hörhistorien in gewöhnliche Musikabspielsysteme zu integrieren.

Diese Dissertation leistet mehrere Beiträge zur Mensch-Maschine-Interaktion und der Informationsvisualisierung: Sie vertieft zum einen das Verständnis relevanter Aspekte und wichtiger Verhaltensabläufe die das Hörverhalten einer Person besonders machen. Weiterhin werden Visualisierungsprototypen sowie ein Designraum für Hörhistorien vorgestellt, der zeigt wie mit persönlichen, zeitgebundenen Daten in einem Lifelogging-Kontext gearbeitet werden kann. Drittens stellt diese Arbeit Wege vor um Empfehlungssysteme und andere existierende Software zu verbessern indem Hörhistorien zur Extraktion der Verhältnisse zwischen Liedern genutzt werden. Zu guter Letzt werden in dieser Arbeit Visualisierungen präsentiert die sich an Laien richten und damit Forschern und Fachleuten Hinweise bieten wie derlei Daten in Zukunft für ihre Erzeuger aufbereitet werden können.

ACKNOWLEDGMENTS

Acknowledgments sections in books, similar to the credits at the end of movies, are necessary but sometimes a little tedious for outsiders. In order to improve the overall experience of reading this section, I would therefore suggest relying on the same trick as film producers and putting on some enjoyable music. An instrumental soundtrack from a war movie should be in order to create the right atmosphere.
First I want to thank my supervisor Andreas Butz for supporting and keeping me on track during the last four years. Whenever I had problems or was not sure how to go on I could rely on him to put things into perspective. I am also grateful to him for being always fair and as supportive as possible for all the main and side projects I embarked on. I am only now able to understand the importance of having a voice of reason keeping the little dramas in perspective.
Second, I cannot thank enough the fantastic colleagues I had at the Media Informatics group in Munich. I was looking forward (almost) every day to being greeted with equal parts of inspiring discussions, fascinating gossip, and utter nonsense. As one of my colleagues, Doris Hausen, once put it: This office feels more like a big shared apartment. I especially would like to thank Sebastian Boring, Michael Sedlmair and Raphael Wimmer for being great researchers and great friends during that time. Every project was enriched by working with them and every conference became much more fun. Having such company makes even a four-year journey feel short. I would also like to thank: Sara Streng (for putting smiles on everyone's faces even while complaining about everything), Alexander De Luca (for miraculously being able to eat and talk at the same time), Alexander Wiethoff (for thinking about people and culture instead of machines and algorithms), Hendrik Richter (for having the driest sense of humor west of Dorfen), Fabian Hennecke (for not taking anything too serious), Ya-Xi Chen (for always being understanding and jovial), Doris Hausen (for being oh so shame-inducingly organized), Alina Hang (for being even quieter than me in the office), Sebastian Löhmann (for not bragging about his friend Don Norman all the time), Max Maurer (for his hands-on approach to the future), and Heiko Drewes (for starting every sentence with a sigh and a shrug and never making good on his threat of emigrating to India). I also want to thank the people who keep this fruitful environment from collapsing, namely Rainer Fink for keeping the machinery running, Anita Szász and later Franziska Schwamb for managing the ever-growing bureaucracy and last but not least Prof. Heinrich Hußmann for always being helpful with answers, suggestions and funding despite all his other duties. A special thanks is due to my former colleagues Otmar Hilliges and Lucia Terrenghi, who endured me during my first year. As Otmar was the one putting me on the road towards research and oversaw both of my theses during my Diplom study this situation, however, was probably deserved.
This thesis would not have been possible without the creative and gifted students that I was able to work with. I would like to thank Frederik Seiffert, Jenny Büttgen, Roman Graebsch and Hans-Peter Dietz for the enthusiasm and hard work they put into their projects. I also want to thank all the other students that I was allowed to accompany

for their theses over the years, most notably Tim Langer, Bartl Steinmayr, and Bernhard Hering.
While working on a project for four years staying sane is the hardest part without people from the outside world showing you that there are other things in life than deadlines and publication lists. I was lucky to have such people. I thank my good friends Max and Basti (for upholding Bavarian culture) and Christoph and Wolfram (for upholding climbing culture) for that. I also want to thank the Werebears (Phil, Dom, Sergej, Nico, Peter) for giving me a reason to learn programming in Objective-C for fun and profit and all the sophisticated discussions about video games.
My biggest thank goes to my family and my girlfriend who were my anchors during these years. To my siblings Benedikt and Elisabeth who had to endure their smart ass big brother. To my parents who always welcomed me with open arms and great food (and cake!) even at inopportune moments. Arriving at the farm far away from the city and the internet always felt like a little vacation.

Finally, I want to thank Bettina for always being there for me, for always understanding my gripes and supporting me at whatever I did. I could not have done it without you and am forever thankful.

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

Introduction

And the band plays some song about
forgetting yourself for a while.
And the piano's this melancholy
soundtrack to her smile.

– The Airborne Toxic Event - *Sometime Around Midnight* –

The ways in which people interact with music has changed dramatically within the last decades. Advances in technology have led to a complete change in abundance and ubiquity of music. Where listening to it was an explicit activity before and required taking time, picking up a physical container (of whatever form) and a suitable playback device, nowadays practically all songs ever written are at the fingertips.

The reasons for this development stem from three main sources of technological progress (cf. [64]): storage formats, network access and device portability. Starting with the inception of recording sound waves on physical data mediums end of the 19th century, music had turned from a fleeting sensation that could only be experienced in a live performance to a repeatable (and replicatable) item. The following development of smaller and smaller mediums with higher data density found its current apex in the change from analog to digital storage and the corresponding forms of applicable compression. The MP3 format effectively increased the available storage space tenfold and flash-based memory chips hold gigabytes of data, leading to an enormous capacity even in portable music players.

The most important aspect of this development was, however, the change from physical artifacts such as records or CDs to purely digital forms. Music had no longer a physical representation or value and became as weightless as all other digital files. Piggyback art forms, such as booklets or sophisticated accessories and private mixtapes created with painstaking accuracy and equipped with extensive drawings were lost in the process. Also, the quality of playback would no longer deteriorate as was the case with vinyl records, cassettes and even digital compact discs. Music could no longer be worn-out and its physical representation would no longer reflect its usage history with scratches and torn booklets.

More advanced forms of compression and digital formats also reduced the bandwidth required for transmitting music. The Napster revolution of the late 1990s set the stage for all (legal and illegal) digital music services that came after it and showed that a direct network access to every computer world-wide would also make the old dream of access to every song imaginable true. And while Napster and other peer-to-peer networks worked with the small bandwidth modems of the early WWW, broadband access to the net made this access instantaneous. The 'jukebox in the sky' - a direct access to every song ever written - has almost become a reality with music subscription and streaming services such as Spotify, rdio or MOG that give access to a selection of millions of items suitable for every passing whim.

Yet, network access has not only revolutionized stationary music listening - fast wireless broadband data networks that have grown out of the voice networks of the past provide equal speeds everywhere. An important addition are suitable portable devices - smartphones and laptops. Miniaturization has led to powerful handheld devices, nowadays mostly operated by touch-screens, that work as interfaces to the cloud-based data providers. While music had already had a headstart regarding portability with devices such as the iconic Walkman and Discman by Sony and the iPod by Apple, smartphones with constant network access are freed from the shackles of local storage and choosing suitable fillings - every song is always available.

These three factors - digital formats, always-on networks and powerful portable devices - have made music what it is today: A permanent part of our lives, accompanying people on trains and in cars, at work or leisure and while alone or in company. Ubiquity has turned it into a constant soundtrack that connects songs written by other people with events of our own lives.

While these technological advances brought more freedom in what to do with and how to listen to music, they also showed the limits of human capacity regarding such media collections. The restricted storage capacities and space requirements of physical media necessarily led to more repetition and made it easier to remember them. Also, they came with colorful covers and booklets that introduced additional properties for finding and memorizing them. Now, even with a portable, non-network player that holds 'only' tens of thousands of songs, keeping track of every single one is close to impossible. And the reduction of booklets to single cover images also took away these additional memory hooks. The greater flexibility that came with the digitization of music therefore introduced new problems of retrieving and remembering. In short: Having millions of songs available very effectively demonstrates the lack of scalability in current interfaces and our own minds.

Fortunately, technology also provides ways to either extend or replace memorization. It is trivial to integrate the capabilities of logging every single song that is played with a playback device. The resulting *music listening histories* provide clear lists of songs that make looking up barely remembered songs and forgotten favorites simple. But beyond that they also give detailed insight into the nature of a person's music consumption: the discovery of new songs or repetition of old ones, daily and weekly rhythms and, by using more than one such history, the social nature of music and the influence of and on other people. In a way, these listening histories bring back some of the aspects of the physical artifact-state of music and reflect how songs were actually used instead of them being timeless, error-corrected digital files. And they even allow doing so in a much more accurate way: Every single listening instance can be tracked and looked up down to the second, something not possible with worn-out booklet corners being all that is left of the past. And with listening histories being digital files themselves, automatic upload to online services and analysis and presentation become possible.

At the moment, a large number of devices and software media players capture these histories, but mostly for reasons that only indirectly benefit the listener: First, listening histories are popular as a source for music recommendation. The abundance of available music not only led to a problem of retrieval but also of discovery. Collaborative filtering [54] is the most popular of algorithms that automatically derive similarity of items from usage or ratings of a large number of people. Listening histories as expressions of approval are commonly used as input to such algorithms for generating direct music recommendations or creating web radio stations. Second, listening histories also inform record companies and other marketing stakeholders about the impact of their products and promotional activities. Listening histories provide a detailed image of the use and popularity of songs in very specific market segments. But again, listeners benefit only indirectly from that with more suitable products. Finally, media player software such as iTunes makes listening histories available but only reduced to the playcount of a song, as a crude metric for the popularity of items in a music collection.

In this dissertation, I argue that these applications only barely show what music listening histories can be used for and that there are many more ways in which listeners can profit from them.

1.1 Motivation

Music listening histories are readily available but their potential has only barely been shown in existing software. As discussed above, even though the listeners themselves produce this type of data, the main stakeholders benefiting from it are the collectors in the form of online services, record companies and marketing agencies. Listening histories promise, however, to provide a much more direct benefit to their creators than music recommendation, either in newly created software tools or as extensions to existing ones.

In a way, listening to music is nothing more than a problem of selecting the right song at

the right time. While there is the overarching set of all music imaginable, only a smaller subset corresponds to the taste of a listener. The fringes of this subset are not stable, as taste might shift from moment to moment depending on the mood or the surroundings, but the core taste only changes slowly if at all.
The problem slightly shifts when the song is part of a longer sequence, as this creates additional restrictions on choice: The next song not only has to fit into the listener's taste but also work especially with the one directly before it and with all the other ones within the sequence [63]. Capturing these listening decisions within a listening history thus also allows deriving some of the context but maybe not the actual reasons.
Also, listening histories are peculiar in that they bridge two main categories of data types (see figure 1.1): On the one hand, they represent aspects of the listener's life and work as a form of memory extension. Similar to other such diary or lifelogging data they allow not only simple retrieval and look-up of items, but also seeing overarching patterns that are harder to find when experiencing them from moment to moment. While there exist various names for such personal, time-based data I will refer to this property of listening histories in the following as their *lifelogging* aspect.
On the other hand, they also reflect aspects of the consumed songs where sequences of them can have specific meanings and display some kind of relation. As listening histories are in this sense mostly independent from their creators and the accuracy of such analyses might even increase with more aggregated histories, I will call this aspect throughout the dissertation the *relational* nature regarding music.
The motivation of this work therefore stems partially from the ease with which this data type can be collected and their direct availability from various services. But the main motivation are the imaginable benefits that integrating these listening histories into various contexts promises for their creators and that only barely have been explored so far.

1.1.1 Listening histories as lifelogging data

As we have seen before, the instant access and practically limitless size of current music collections creates a problem of scale for retrieval and remembering. Music listening histories can capture a person's interaction with music and provide similar benefits to other lifelogging systems (cf. [149]): First, they simply help with finding forgotten items. Apart from this *retrieval*, they also work for *analysis* and exploring patterns on a higher level than single items. Finally, as a more sentimental use, they also allow *reminiscing* about the past, similar to photo albums.
One major problem with using music listening histories in their lifelogging sense is that they can become extensive and complex: With an average song length of several minutes a single week can already hold over 1,000, possibly repeating, items. As all songs are uniquely identifiable direct textual search is possible, but not helpful when title and artist are forgotten. Similarly, reminiscing about past experiences is difficult with only long lists of song titles and dates available. To create working experiences with this data type, the right presentation becomes central. As with all large data sets, utilizing

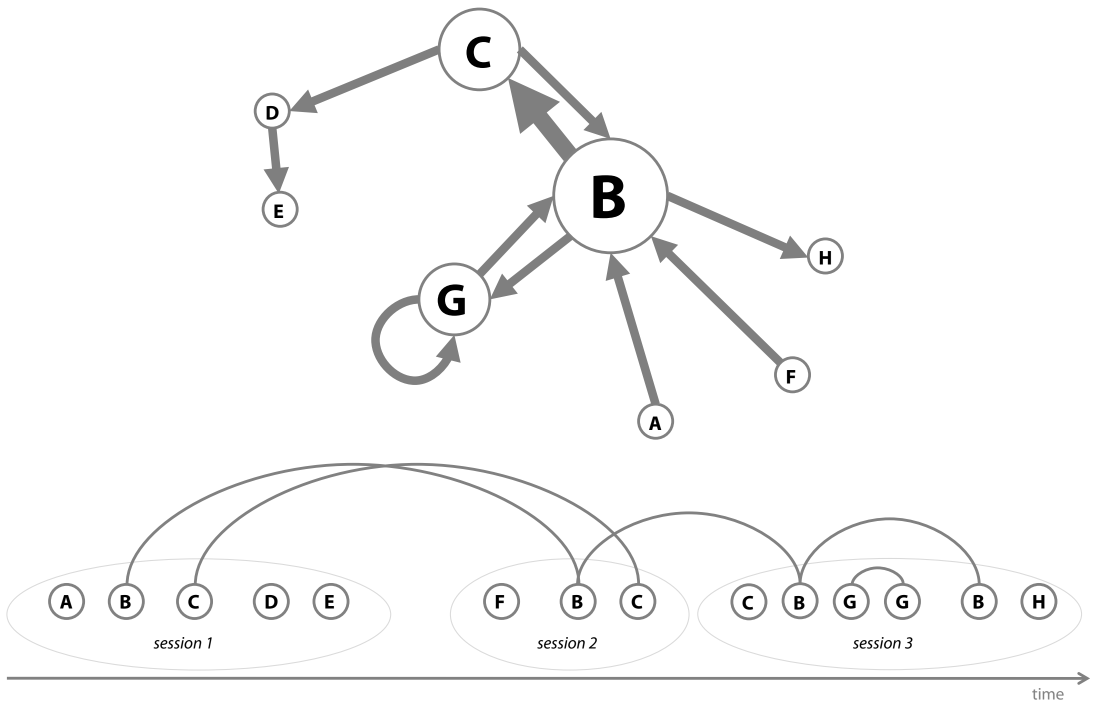

Figure 1.1 Music listening histories can be interpreted in two ways: First, they represent personal temporal data that logs the listener's life. Second, they can also be seen from the perspective of the musical items, with related songs being in the same session and popular one's repeated.

the large bandwidth of the human visual system enables more insight than textual or automatic ways. I therefore chose information visualization as a set of techniques to present and work with listening histories in this sense. Also, as the creators of this type of data are not necessarily experts with visualizations, keeping the resulting tools straightforward and task-centric (*casual*) was important.

1.1.2 Listening histories as relational music data

Beyond the personal meaning ingrained in listening histories, they also reflect attributes of the consumed items. Songs are only rarely listened to in singular instances but come in sequences: most commonly as predefined albums, in an order as intended by the artists, but also as manually and personally created playlists or even in an automatically generated fashion from web radios or playlist generators. As people often choose these sequences deliberately (as they have to listen to them after all) or at least skip songs they do not like [125], items that appear in groups within listening histories share a certain meaning - the closer they are in this sequence, the more expressive. When looking at a

single listening history these sequences represent mostly personal taste, but they become more expressive in an aggregated form. A large number of histories removes quirks and represents a more overarching identification of song relationships. The aforementioned collaborative filtering also creates recommendation based on this principle.
Based on these two extremes of a single history versus a large number of them, sequence-based similarity can be used to enhance a listener's interaction with music, either for her or his own taste or on a general level. Yet, while the underlying data is still complex, using information visualization tools distracts from the actual tasks, such as creating a playlist or just listening to music. Integrating visualization concepts into the systems, however, can help listeners benefit from this data without standing in the way.

1.2 Research Objectives

In order to structure my work and explore the possibilities of integrating music listening histories as representations of lifelogging or relational data into software tools I formulated two main research questions:

Research Question 1: What patterns can we find in listening histories?
As a first step towards using listening histories, we have to understand what this data type contained and what salient features distinguished one history from the next. On the one hand, this question boils down to the ways people listen to music: reasons for choosing a song or a sequence of songs are manifold and might stem from various sources. The current environment and listening device is just as important as the mood or even the task that the listener wants to accomplish. Also, external factors such as the company and restrictions such as a library limited by the storage capacity of a mobile device can influence these decisions and lead to compromises.
On the other hand, the expressiveness of listening histories also depends on how much of the data can be captured. Real world tracking of such behavior necessarily fights with inaccuracies and failures which can come from various sources. Distinguishing the peculiarities of one listening history compared to another one also requires taking noise within the data into account.
As mentioned above, listening histories are also interesting from the musical, relational perspective. Knowing where such sequences come from and what the underlying logic might be can also help in understanding and expecting patterns in the data type. Understanding the musical medium and concepts such as genre or types of sequences such as albums or playlists are necessary when presenting them to listeners.
The result of answering this first research question gives us a good overview of what to expect from musical listening histories as a data type and forms a prerequisite for the second part of actually doing something with them.

Research Question 2: How can we provide people with benefits based on their listening data?
Once the characteristics of the data type of music listening histories are known, we can derive applications for the data. This implies getting an overview of how listening histories are used right now and where to get them from. As RQ1 is also about the ways people listen to music and their needs in doing so, we can then think about the best way to support them in these tasks with listening histories. Having an overview of music tasks allows us to create a mapping between the properties of the data and the possible applications.
As mentioned above, listening histories are personal as well as relational data and possible applications will most likely fall into one of these categories: either as vehicles for nostalgia and analysis of personal behavior, or supporting playlisting and working with music. But we can also expect to have one type of tasks inspire the other one: even though the focus of a tool might lie on analysis, it can also be a nice touch to be able to play back a song from within the interface.
Also, possible applications depend on the amount of listening history data available for integration: One listener can benefit from her or his own history, but having more than this one available allows putting it into perspective or deriving more relational aspects of the data. In the most extreme case, having millions of listening histories available allows creating recommendations or determining "overall" similarity between songs.
This second research question will provide practitioners with new ideas on what to do with this type of data and might also create a need in listeners to demand the intergration of listening histories in their existing tools (especially as the data is mostly readily available).

1.3 Research Approach

The two research questions of how people listen to music and what to do with the resulting data already imply a sequential approach to answering them. Determining how to integrate listening histories into user interfaces can happen after discovering every single new aspect, but having a bigger picture of the nature of this data type leads to a more structured approach. Therefore, I divided my research into three specific tasks to answer the research questions.

1.3.1 Identify people's listening behavior

As a way of answering RQ1, I performed a literature review to identify salient features in listening histories. With music being an important topic throughout human history, various disciplines concern themselves with the impact of music on societies and individuals. Scattered throughout this research and worded in different ways are the factors that influence listening decisions. For the first part of my work, I wanted to get an overview of what these factors are in order to be able to look for them. While this list of factors incorporates all relevant aspects that influence listening decisions, deriving a model of human music listening would have gone far beyond the scope of this thesis: As most observations about listening in research have happened either on the high level of musical genres or for small groups of listeners, researchers are still at the beginning of creating a model of human listening (for an exemplary approach see [178]). Approaches to large-scale analyses of such data sets (see section 3.4) promise to shed more light on human behavior in this regard and provide lower-level descriptions of listening choices. For integrating listening histories with user interfaces, however, just knowing these factors would suffice to distinguish between more or less interesting aspects of a history and also present them in the right way.

1.3.2 Retrieve and characterize real-world listening histories

Using listening histories requires not only identifying relevant listening factors but also acquiring suitable data. While there are several companies and services available right now that collect this data, only a few allow accessing it in the detailed form required to reflect the low-level factors.

Besides finding a suitable data source, it was also necessary to determine which of the factors were actually available in the data: The literature review uncovered the ideal version of all factors that influence listening decisions. Real-world listening histories can naturally only display a subset of these factors, as their approach is often strictly focussed on the music without all the necessary contextual information. Also, many of the factors are dependent on the listener's current mood and other hardly measurable aspects that just cannot be automatically captured in the data. I went on looking for other data sources that could at least contain some of this information. Especially the listener's memories can be a valuable source for this rather hard-to-grasp data. Triggering these memories is a necessary prerequisite for making use of them, but was at least possible with the available data.

Besides creating a matching between listening factors and listening histories, I also analyzed real-world histories with regard to noise and general problems in the data. Various factors influence the accuracy of listening histories and problems arise mostly from too little songs tracked but also from erroneous trackings. Presenting the data requires being aware of these problems to help listeners understand and find such problematic sections. Also, the amount of errors in the data was important, as mostly false histories are practically unusable.

Finally, having actual listening history data available also allowed me to provide estimates about the impact and importance of listening factors. Analyzing a large number of histories made it possible to see which factors were most expressive in determining differences between two listeners and the impact of aspects such as demographics on music listening.

1.3.3 Integrate the data with user interfaces

With the data and its characterization available, I could in a last step create a mapping between music tasks and the listening histories and answer RQ2. In this main part of my work, I derived a design space from the listening factors and the intended tasks surrounding music. Based on this design space, I created several prototypes in an experimental approach.
These prototypes were on the one hand used to explore the data and provide insight into its structure, as listening histories with thousands of songs already need specific tools to work with. On the other hand, the prototypes worked as exemplary implementations for using listening histories in a lifelogging or music context. As the underlying data is always extensive and complex independent from the use case, all prototypes rely on information visualization techniques for creating insights or handling the task, but they do not necessarily resemble "classical" information visualization tools. Where appropriate, I performed evaluations of the implemented prototypes, either after implementation (summative evaluations) or to learn more about people's needs with their listening histories beforehand (formative evaluations).
The first set of applications (see chapter 5) emphasizes the lifelogging aspects and have analysis and reminiscing as main tasks. While these two distinct tasks encourage using different interfaces, I tried to make it as easy as possible to switch between the two. Also, depending on the question that can be answered with the prototype, complexity in the interface can be reduced.
The second set of visualizations (see chapter 6) shed the analytical looks of the former prototypes and work in more music-centric contexts. They focus on the problems of actually listening to music, but also doing reminiscing or rediscovery while listening to it. These prototypes use the musical relationships expressed through listening histories for these tasks and present them in ways inspired by information visualizations but without the complexity.
The design space and the prototypes as well as the corresponding evaluations show the potential of integrating listening histories with music applications.

1.4 Contributions

This thesis provides several contributions to the fields of information visualization and human-computer interaction. It also points out new ways to approach learning about human music consumption. Finally, it also provides practitioners with new approaches to presenting information to regular people and integrating consumption histories. In this section I will present these contributions in a more abstract way, while coming back to them at the end (see section 7.1) to show them in light of the presented work.

1.4.1 Deepening the understanding of human behavior surrounding music

As a first contribution, this thesis merges and structures insight into listening decisions and factors. While a large amount of research has described various aspects of this part of human behavior, their results are scattered throughout musicology and computer science. In this dissertation, I provide a list of factors based on this research that are relevant when thinking about why a person listens to what music. This list of factors is useful for researchers as a checklist for possible questions and observations as well as for practitioners that want to have an overview of relevant aspects in their music tools.
Further, I also show how using real-world tracking data can be used to improve our understanding of human listening: taking a large sample of listeners and looking for patterns instead of formulating hypotheses and validating them using a small sample [4] uncovers more independent and specific insights that can also be less obvious than the ones acquired manually. Ever-growing sensor networks and the inclination to share more and more data online makes this approach to learning about human behavior promising for science.

1.4.2 Structuring the design space and suggesting ways for visualizing personal temporal data

The inclination to share more data with online services also implies the need to make use of it. As a second contribution, I provide a design space for prototypes that work with listening histories or other related data (mainly other consumption histories). This design space structures existing and future prototypes and shows areas where clear solutions exist and areas where they are needed. It can also be used in initial stages to come clean with what a prototype is supposed to do.
In this dissertation I also present various exemplary prototypes that are parts of this design space. While they are specifically set on music listening histories, they, just as the design space, can also be used as inspiration for other applications based on related data. Also, the subsequent online evaluations showed what approaches were promising, what tasks were useful to people and where the sweet spot between functionality and complexity lies.
Finally, these prototypes also provide insight into what techniques are effective for non-infovis-experts. I also addressed the problems of enticing people to use something and providing them with easy ways for learning.

1.4.3 Using visualization approaches for music-related tasks

As a last contribution, this dissertation shows how concepts and techniques from information visualization can be used for other goals than to provide insight [29]. For complex data such as music or other media collections, visualizations can be included throughout the interface to inform about underlying structures or even work for direct manipulation of abstract data. Especially the second, music-related set of prototypes (see chapter 6) shows how concepts from information visualization can work seamlessly for regular tasks such as building playlists or listening to music.

With the second set of prototypes I also show that these concepts can be integrated into usage contexts different from regular desktop-based setups: Changing from the desktop to a mobile device or integrating listening histories into existing media player software (instead of creating another standalone software tool) can make more sense, depending on the intended task. Also, relying on software that the listeners would have used anyway for their purposes eases the amount of required learning and lowers the adaptation threshold. Newly created software can make use of these concepts and integrate listening histories from the start.

Finally, this thesis also contains a more detailed concept of similarity based on co-occurrence within histories. While concepts such as collaborative filtering very efficiently make use of this automatically generated similarity-metric, providing listeners with their very own, non-aggregated versions of these listening sessions and predecessors and successors provides its own benefit.

1.5 Thesis structure

This thesis is structured as follows:

- *Chapter 1 - Introduction* provides an introduction to the topic of music listening histories. It also discusses the motivation behind this thesis, of using this easily collectable type of data for their creators' benefits. In this chapter I also present my research questions and the approach that I chose to answer them.

- *Chapter 2 - Human-Music Interaction* contains an extensive literature review from musicology and computer science about ways in which and reasons why people listen to music. In this chapter, I also give an overview of listening factors that influence and restrict musical choices.

- *Chapter 3 - Music Listening Histories* details the data type of music listening histories. I present related work from lifelogging research and discuss listening histories as another form of personal, temporal data. After that, I describe the properties of actual listening histories and present the results of a large-scale analysis of such real-world histories. The chapter closes with a discussion of the mapping between the listening factors of chapter 2 and the actual data.

- *Chapter 4 - Visualizing Listening Histories* is the basis for the next two chapters and provides background on the task of visualizing this data type. Apart from discussing other exemplary visualizations, I derive a design space that captures the most important features of listening histories.
- *Chapter 5 - Listening Histories: Visualizations for analysis and reminiscing* gives several examples for listening history visualization for the specific tasks of analysis and reminiscing. This chapter emphasizes the lifelogging aspect of such histories and presents more "classical" visualization approaches. Each description of a prototype closes with a short discussion on the main outcomes of the respective work.
- *Chapter 6 - Listening Histories: Visualizations for playlisting and rediscovery* provides a second set of prototypes with the focus on music-related tasks and non-desktop scenarios. Another central topic of this chapter is the relationship between musical items expressed in listening histories (either in accumulated or singular form). Each prototype also closes with a discussion part.
- *Chapter 7 - Summary and Future Work* closes the dissertation. After a more in-depth discussion of the contributions, I talk about the impact of personal visualizations in the future.

Chapter 2

Human-Music Interaction

He's the one
Who likes all our pretty songs.

– Nirvana - *In Bloom* –

The role of music in society has changed enormously throughout history. While music is with our species since the onset of civilization, its use was initially coined by the necessity to perform a musical piece live due to the lack of ways to record it. Therefore, music became something special and was no everyday commodity: Important events such as holidays were accompanied with live performances and the impact of religious ceremonies was enhanced through the addition of the rarely-heard music. Even in relatively modern times before the 20th century music was "a heightening of the everyday into the special" [142] and visiting the opera or a concert were important social occasions. Kings and other rulers had their own orchestras at their courts and gloated with their composers and the rising bourgeoisie took pride in having their children learn instruments.

All that changed with the invention of the phonograph and the subsequent recording of music. Suddenly, music had become something that could be owned and enjoyed by oneself without the need for musicians present. Also, each recording only needed a single performance and nothing more - so actually living of performing music became harder. As North et al. express it: "Music can now be seen as a resource rather than merely as a commodity" ([120], p. 42).

But this freedom also brought additional responsibilities for music listeners: While it was and still is possible with live musicians to just let them choose what music they play, now the listener had the new tasks of acquiring new music on whatever medium, organizing this collection of either analog or digital music artifacts and, most important of all, decide whatever song they were currently in the mood for. The listener was no longer a passive participant - as the name implies - but became at the same time the director, technician and (hopefully) entertained audience. In addition, the task of curating a collection also fell to her or him. In essence, the new-found freedom held also the responsibility to keep oneself entertained, with a failure being no longer some musicians' fault but one's very own.

In this chapter I first give an overview of what reasons people have for listening to music and what tasks they have to perform with their music collections. I then present several studies from musicology and human-computer interaction about the ways people listen to music. The chapter concludes with a list of listening factors that determine when people listen to what music.

2.1 Uses of Music

The main question that all scientific disciplines concerned with music try to answer is the extensive use of music by the human race. Especially today, music seems to be everywhere in various forms and contexts. Levitin and McGill call it 'life soundtracks': "music we listen to that inspires, motivates, calms, excites, and generally moves along the action in our daily lives" [102].
Yet, not all music is consumed deliberately - in various contexts (e.g., dining areas [115], aerobics classes [44], while shopping [110]) music is added from external sources and other-directed. Oftentimes, the reasons for adding music to public locations and forcing others to listen to it are commercially motivated (several studies on consumer behaviour have shown the effectiveness of music for adjusting a customer's perception of time ([177]) or shaping product preferences ([121])), sometimes music is used for creating a communal feeling (especially in religious or public celebrations), and sometimes it is just there due to negligence or even malevolence (e.g., the proverbial kid with the boombox).
For the sake of this discussion, however, I assume the listener to be completely self-determined. She or he is the only one responsible for choosing the music that is playing and decides whether to listen to music at all. Sometimes, there might be other people present and in these cases the listener's choice can be influenced by this social context (see below). Discussing personal uses of music seems more suitable when all relevant choices are made by the listener.
Also, while the central performance aspect of music dwindled throughout the 20th century, new electronic ways of creating and remixing music on-the-fly have led to a resurgence of the consumer as creator. Especially with the help of content- and context-aware tools (e.g., RJDJ [1]) music listening in the future promises to provide a

[1] http://rjdj.me/

unique experience every time a song is played. For this chapter, however, I presume music to be static and the listener only being able to choose from it and turn it on or off.

A common approach to explaining the use of music by people is assuming 'uses and gratifications', a theory originating from research on mass communication ([96]) but applicable to various other modes of consumption. Contrary to earlier approaches of explaining mass communication, Katz and Foulkes asked the question "not 'What do the media do to people?' but, rather, 'What do people do with the media?'" [80]. Media consumption was no longer controlled by the content producers and distributors but by the consumers themselves, a view that might have also stemmed from the beginning introduction of choice into mass media, with more and more radio and TV stations available at the time. This emphasis on choice is in line with the individualism and personal taste especially relevant to music.
The central idea of uses and gratification is that any (conscious) *use* of a medium is motivated by some proposed *gratification* sought after by the consumer. The listener has some explicit reason for choosing to listen to music and picking a certain song. The theory assumes a listener that expects a direct causality between media consumption and some goal. In a way, the idea of man of uses and gratification is similar to the *homo oeconomicus* of economics: A person is a perfectly rational actor, at any point in time completely aware of her or his wishes and ways to fulfill them. While this perception of human existence might be debatable, choosing it for explaining the use of music gives at least a rough idea of what is happening behind the scenes. Even though people might not obtain the gratifications they sought after with a certain music choice (especially as we are notoriously bad at estimating the impact of choices on our happiness [98]), they still have at least a motivation and an idea of what they wanted in mind. In the following, all studies that ask for reasons, for the *why* of a music choice, are implicitly building on uses and gratifications.

One of the first categorizations geared towards understanding the human use of music comes from Merriam [109], who names ten reasons for listening to it: "emotional expression, aesthetic enjoyment, entertainment, communication, symbolic representation, physical response, enforcing conformity to social norms, validating social institutions and religious rituals, the continuity and stability of culture, the integration of society" (cited after [64]). These uses not only contain individual reasons, but describe more generally also societal reasons (e.g., "validating social institutions and religious rituals") that might lie outside of the private choice and fall into the other-directed category. Also, it contains reasons that people not consuming but creating music might have: Merriam mostly talks about "emotional expression" and "communication" in the sense of self-expression through music, with music being a vehicle for art and emotion (even though it can also be interpreted in the sense of expressing yourself via a music choice).

Various other approaches to determining and classifying the use of music exist in music psychology. North et al. asked their participants in a large-scale study [120] very generally what effects the music had on them and provided a list of eleven factors: "It helped me concentrate/think, It helped to pass the time, It helped to create the right atmosphere, It brought back certain memories, Someone else I was with liked it, Helped create or accentuate an emotion, Helped create an 'image' for me, Habit, I wanted to learn more about the music, I enjoyed it, Other". While these factors are certainly important, they are (deliberately) kept redundant (e.g., 'It helped to pass the time', 'I enjoyed it' and 'Habit' all seem related) which hints at a more general classification. Roe used a similar approach in his study with 500 Swedish teenagers [140] and his possible answers to the uses of music are not very different from North et al.'s: "It helps me to relax and stop thinking about things, It helps me to get into the right mood, It helps pass the time, It is less boring when I am doing something else, It is good to dance to, I feel less lonely when I am by myself, I want to listen to the words, It fills the silence when no one is talking, It creates a good atmosphere when I am with others, Music fits in well with my life, The words express how I am feeling, It makes the time go faster when there is nothing to do". Roe provides a more general classification for these answers as *atmosphere creation and mood control*, *silence filling and passing the time* and *attention to lyrics*, but while this is helpful, it only makes up a small section of what uses of music can mean.
A very basic but general approach comes from Rösing: "When, therefore, music exerts an effect on, and 'functions' on, the receiver, it does so in at least two respects: in the communicational/social sphere, and in the individual/psychological sphere." [142]. Roughly spoken, music either has an inward effect, on the listener her- or himself and an outward effect on the social surroundings.
A threefold approach to classifying the uses of music comes from Hargreaves and North [64] and Chamorro-Premuzic et al [32, 31]. They both see three main categories: *cognitive / intellectual / rational use*, *emotional use*, and *social / background use*. Especially Hargreaves and North argue that while the first two uses of music have been addressed extensively in research, the social aspect has only barely been scratched (they also do not miss the chance to mention their own book on the 'The social psychology of music' [66] in this regard). Chamorro-Premuzic et al. name the three uses as a basis for examining the connections between personality traits and listening behaviour [31, 32] (see below for a more extensive discussion of these studies). The threefold approach of these authors adds another dimension by separating Rösing's 'individual sphere' into an emotional and intellectual section.
Rentfrow and Gosling were the first to systematically examine the connections between personality and music taste in 2003 [136] and generally agree with the threefold categorization. They, however, emphasize the interesting aspect of 'identity claims' via music as a bridge between social (other-directed identity: "individuals might select styles of music that allow them to send a message to others about who they are or how they like to be seen" [ibid.]) and emotional (self-directed identity: "individuals might select styles of music that reinforce their self-views" [ibid.]) needs.

Based on these classifications, I will describe the three main uses of music as *mood management*, *cognitive functions* and *identity construction*.

2.1.1 Mood Management

The emotional uses of music can be subsumed under the term 'mood management'. Initially coined for TV use, Zillmann adapted the theory for any media use [180]. As Zillmann says: "The theory posits that individuals are capable of choosing materials for exposure that modify and regulate affective experiences and mood states in desirable ways, and that these individuals frequently and habitually make choices that actually serve the specific ends" [180]. Mood management again assumes along the lines of uses and gratification that the listener is able to estimate and select media for positive modifications of their current emotional states.

In a classical study [87], Konečni et al. demonstrated the "arousal moderation" part of music: Based on existing psychological research about the influence of complex stimuli on arousal, they hypothesized that an angered participant of their study would tend towards a lower complexity stimulus to manage this arousal. They separated their participants into three groups: The first group was angered but could punish the aggressor using fake electroshocks, the second was angered without any way to relief the tension and the third was the unangered control. Afterwards, all participants could choose between a simple and a more complex melodic sequence to listen to. Participants from the second group significantly more often chose the simple melody, in order not to increase their existing arousal. Konečni et al. also performed other creative studies that showed the impact of melodic complexity, listening volume and social interaction on mood (for an overview see [86]).
Beyond these basic indicators, mood management also affects more complex musical structures. Knobloch and Zillmann showed in a controlled study [84] how popular music is selected for mood management. Participants first completed a test and the experimental group was intentionally frustrated by using negative results. Afterwards, participants were asked to pick their favourite songs from an available selection, while listening duration and selection were recorded. The results showed that "respondents in a bad mood preferred exposure to highly energetic-joyful music over music low in these qualities to a higher degree than did respondents in a good mood" [84]. Conversely, participants in a good mood listened to significantly more different songs in the available ten-minute period.
Even though the results of these studies show the impact of music on mood management, North et al. criticize the approach: "However, the ecological validity of such experimental research may be limited as a result of its experimenter-centered approach. Participants were made to listen to music of the experimenter's choosing in a situation of the experimenter's choosing and to select from several response options of the exper-

imenter's choosing" [120]. Their argument is reflected in ethnographical studies, where options for reasons for listening to music commonly contain thinly veiled versions of mood management: "Helped create or accentuate an emotion" and "It helped to pass the time" in [120] or "It helps me to get into the right mood" and "It is less boring when I am doing something else" in [140]. These options being a popular choice in the respective studies also showed the applicability of the mood management theory in real-world conditions.
Finally, Levitin et al. presented very interesting work on the cognitive neuroscience of music (for an overview see [103]) and its actual neurophysiological impact on emotions. Menon et al. showed [108] that listening to pleasant music has a similar effect on the brain's reward center as sexual intercourse or opiates. Grewe et al. examined [61] the phenomenon of 'chills' while listening to music and were able to map strong emotional reactions to music attributes (e.g., the "Barrabam-call" from Bach's Matthäus Passion [61]).

2.1.2 Cognitive Functions

A peculiar aspect of mood management is the effect of the applied stimulus on a person. The intensity of this effect cannot be clearly established (except in very abstracted situations such as Konečni's simple melodies [87]) as people have different tastes in music and different personalities. Also, while the experiments above were mostly based on arousal-aversion, mood management and the related Reversal theory [6] hypothesize that people with a low level of arousal will seek arousal-laden stimuli against their boredom. Music can provide different levels of arousal depending on the listener's capabilities: Complex music might be just right for someone with higher cognitive abilities, but can be overwhelming for somebody with lower ones.
In their study from 2003, Rentfrow and Gosling [136] compared music preferences and various personality attributes. As a first step, they separated the 14 most common and popular music genres into four main psychological dimensions: *Reflective & Complex*, *Intense & Rebellious*, *Upbeat & Conventional* and *Energetic & Rhythmic*. These dimensions were extracted with a principal component analysis of the ratings of 1.700 students and turned out to be the most salient factors. In a follow-up-study, the authors examined music preferences, personality traits (using the Big Five Inventory test [75] amongst others), cognitive abilities and self-views regarding politics, etc. The results showed various connections between personality and music preference: One result was that "the Reflective and Complex dimension was positively related to Openness to New Experiences, self-perceived intelligence, verbal (but not analytic) ability, and political liberalism and negatively related to social dominance orientation and athleticism" [136]. The opposite dimension of *Upbeat & Conventional* also had an inverse set of correlated attributes. Apart from the four dimensions, the authors also mention the "the absence of substantial correlations between the music-preference dimensions and Emotional Stability, depression, and self-esteem, suggesting that chronic emotional states do not have a strong

effect on music preferences" [ibid].
Chamorro-Premuzic et al. followed this study with two studies of their own [31, 32], again about the impact of personal traits on music listening. They worked with similar personality measures but did ask participants not only to rate music but additionally provided them with a questionnaire about their uses of music. In the first study ([31]) they worked with British and American students and switched to Malaysian students in the second study ([32]), to control for cultural effects.
They analyzed the questions on use with a PCA, arrived at three main uses of music (emotional, cognitive and background) and separated the last one into pure background and music in social situations. The mapping between personality factors and these uses showed unsurprisingly that persons with high values of neuroticism and introversion had a tendency towards emotional uses of music, while cognitive use was closer correlated with higher cognitive abilities and a disposition for intellectual tasks. Their second study confirmed the main results of the first. In general, both studies were also consistent with Rentfrow and Gosling's study above, especially regarding intellectual attributes and preference for complex music.
A standard rational use of music is for supporting other tasks, either for making tedious tasks more enjoyable (e.g., housework, riding the bus) or improve the performance in exercise. While the former use relates to the low arousal of the main task and the mood management theory, the latter can improve the performance - or at least the enjoyment - of exercising. Vignoli [167] and Cunningham et al. [41] describe independently from each other musical categories such as 'programming music', 'driving music', 'work music' and 'music to amuse children' among others in their studies on music organization.
Levitin and McGill [102] also mention the theory of flow with regard to music. Csikszentmihalyi's idea of flow [36] states that people aspire to an optimal experience by finding the middle ground between too simple and too complex stimuli. These stimuli can of course also be musical in nature to help in performing "at peak mental and physical capacity" [102].

2.1.3 Identity Construction

Music also plays a central role in social activities [64]. Apart from accompanying ceremonies and celebrations, it also helps in defining one's own identity.
The participants in one of Rentfrow and Gosling's studies "certainly believed that knowing a person's music preferences could reveal valid information about what he or she is like" [136] and the connections between personality and preferences regarding music hint at the possibility of deriving personality traits from a person's music taste. In a later study, Rentfrow and Gosling examined that hypothesis [137]: Participants performed several personality tests and compiled a CD of their top-ten favourite songs. Judges afterwards had to estimate the participants' personality traits on the basis of the CDs alone. The results were different from other such "zero-acquaintance" studies (e.g., with photos or the participants' rooms [56]), but showed "that individuals' music preferences

convey consistent and accurate messages about their personalities. Additionally, the results suggest that specific attributes of individuals' music preferences and music-genre stereotypes differentially influenced observers' impressions of targets' traits, values, and affect" [137].
The authors also distinguish between self- and other-directed forms of identity management. Gosling describes the concept of self-directed identity as follows in his book 'Snoop': "It is a story you tell about yourself to make sense out of what has happened in the past and the kind of person you are now. From this perspective, it is not essential that the story be true" [56]. Accordingly, music that one listens to and is aware of can help reinforcing this idea of identity. This confirmation is another emotional use of music.
Another aspect of this construction of self-identity is reminiscing. In remembering and re-telling one's story, self-identity is reinforced and one's own memories and behavior appear more grounded and justified. Similar to photos, music is a popular media for reminiscing. Bentley et al.'s present several exemplary uses of music for reminiscing in their study on the similarities between personal use of photos and music [16]. Based on these results, the authors suggest building a system which "could automatically create a CD for a particular vacation or event for a particular person based on music most frequently played at that event or with that person" [ibid.], similar to Sellen and Whittaker's *reminiscing* use for other lifelogging data [149].
Other-directed forms of identity management have become more important in recent years with the growth of social media: Much of the work that I did in the context of this dissertation was with public listening histories from the last.fm webservice that presents an automatic platform for sharing one's listening habits. Especially this automated sharing proves to be problematic in conjunction with upholding a desired self-image online. Silfverberg et al. [153] talked to several last.fm profile owners about their opinions on and strategies for managing their profiles. While people had developed various strategies for adjusting their profiles ("Playing one 'wrong' song and diluting it by playing several 'good' songs", "Switching the scrobbler off and listening to a song in private"), they even went as far as 'boosting' a song by playing it without actually listening to it. Interestingly, participants thought that "Interpretations, whether negative or positive, were considered unfair and against the prevailing trend of trying not to be judgmental" [ibid.], contrary to the idea of using the profile as an identity stand-in.
A less intrusive form of online self-representation was examined by Voida et al. [169] in the form of shared iTunes music collections and playlists. While the identity presentation was less detailed, the environment of the internal network of a company and the professional relationships of the participants added other constraints (one participant noted: "Some people have expressed some concerns especially when the managers started sharing, started browsing other people's collections"). Participants in this study certainly were also aware of the implications and their impressions on others.
Voida et al. also mention the concept of 'playlistism', when self-identity is (unintentionally) established through publicized playlists. Invented on college campuses, it gave

students a way to "[discriminate] based not on race, sex, or religion, but on someone's terrible taste in music" [79].

These three main aspects, *mood management*, *cognitive functions* and *identity construction* provide a good overview of why music is used in everyday life. These uses, however, are mostly independent from the underlying technology used for listening. While all authors are aware that "[t]he fact that music, by means of the technological transmission-chain, is available at all times and can be replayed at will, independent of the here-and-now of live performance, has far-reaching consequences for the listening behaviour, listening expectations, musical preferences and musical understanding of every individual" [142], exploring the actual technological aspects are not the musicologists' central goals. What these music tasks constitute has been studied more extensively by people responsible for the technology-side.

2.2 Music Tasks

The new listening technology has not only contributed to a slow decline of live performances but also changed the role of the passive listener to that of an active entertainer and curator. Music artifacts, be they digital or physical, can be collected and must be organized and accessed to be useful. The formerly trivial act of listening to music has been subdivided into more of a librarian's job than that of a connaisseur.
Tasks surrounding music can be separated into a chain from discovering a piece of music to consuming it (see figure 2.1): After *discovering* a song it has to be *acquired* and included into a music collection that has to be *organized* and where *understanding* can happen. Additionally, favorite songs might be *shared*, but the final goal is always *listening*. For each section of this chain exist various software tools and some of them (e.g., Apple iTunes) even attempt to support all of the steps.
These tasks are repeatedly mentioned - mostly implicitly - throughout the literature (cf. "music lifecycle" [22]) and there exist dedicated studies for each step of the way (discovery: [39], acquisition: [23], organization: [41, 16], understanding: [8] , sharing: [146], listening: [101]).

2.2.1 Discovery

With the enormous amounts of music available today (most music streaming services boast more than ten millions of songs [2]) actually finding what fits one's mood or intent can become complicated, just as the opposite case of not wanting to miss personally interesting music.

[2] http://pansentient.com/2011/03/music-streaming-services-who-has-the-most-music/

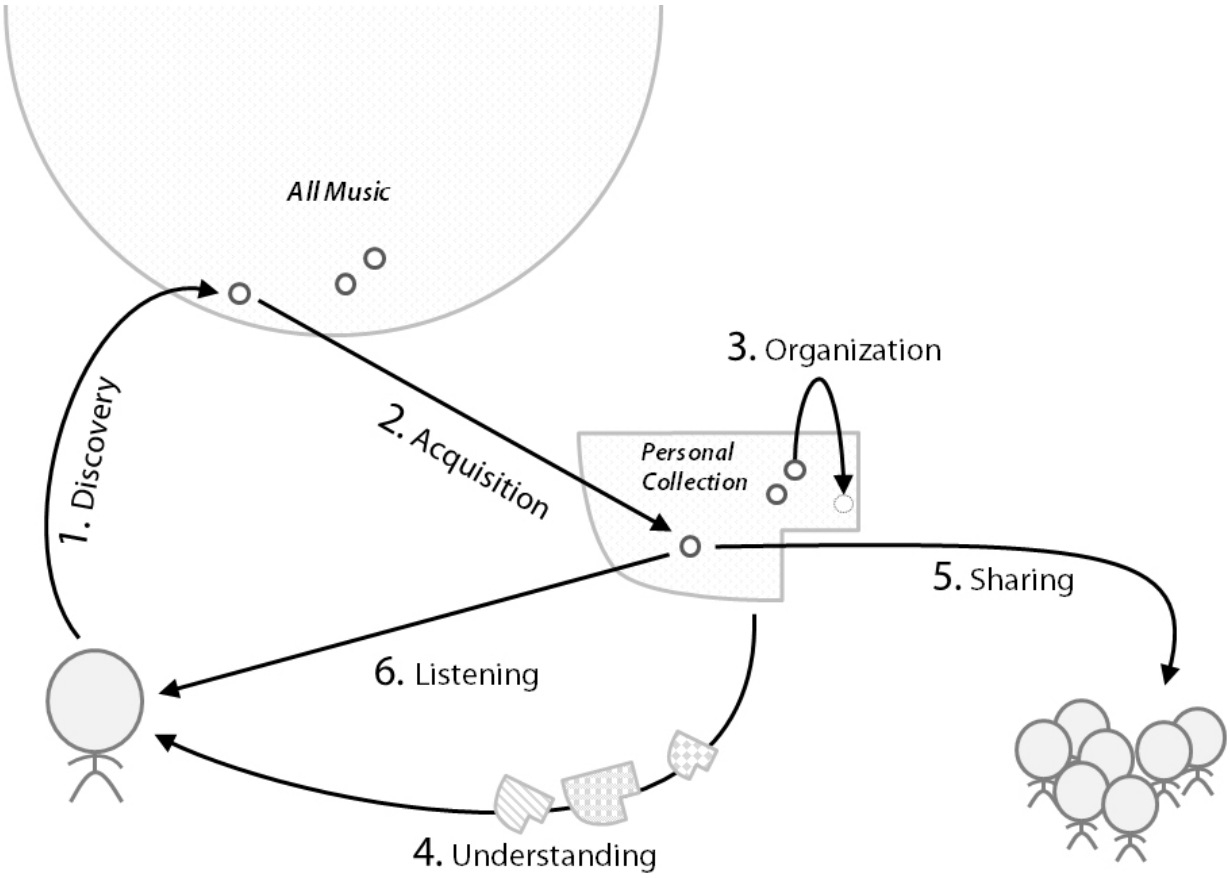

Figure 2.1 Interaction with music happens in six steps: Songs are discovered (1) and then acquired (2) and added to the personal collection. This collection is subject to organization (3) and a listener's understanding (4) of it. Songs are shared (5) with others and also listened to (6).

Consumers follow several different strategies for discovering relevant music. Cunningham et al. followed music shoppers in an ethnographic approach [42] and describe, for example, the "monitoring information seeking strategy: [customers] peruse the store displays just to stay current on the locally obtainable offerings". With the 'extraction' strategy "they are [working] systematically through a particular group of CDs in order to identify material of interest, much as researchers engaged in a literature search may examine a series of issues from a journal one by one in hopes of locating a useful research article" [ibid]. Cunningham et al. also emphasize the importance of the visual impression while searching for music.
Apart from physical CDs that enhance browsing behavior and exploratory searching, media have always played an important role in discovering new and interesting music. While the culture of music journalism traditionally lived in explicit music-centric journals or as reviews in regular newspapers, it is increasingly shifting towards the web as a platform [74]. Jennings follows the idea that with the web as democratizing force

amateur journalism in the form of private blogs is becoming more and more important, especially as blogs have an advantage in trust compared to traditional media (cf. [76]): "[Reading a good blog] has the personal touch that lends it an immediacy and authenticity we don't get from more mass-scale professional media" ([74], p. 146). In this regard, both online and offline media are sources of inspiration for finding new music. Yet, the problem of scalability also became relevant for the sources themselves: The amount of music blogs and other related publications has increased similarly to the treated topic, which makes automated approaches sometimes more feasible.

Beyond unknown music, the ever-growing music collections and unlimited streaming services have led to the interesting phenomenon that by now music discovery is even possible within one's own music collection. Oftentimes, people just download and hoard as much music as they can or have streaming access to tens of millions of songs so that intimately knowing each of these items is no longer possible. Therefore, playlist generators such as iTunes' Genius can also be used for discovery within a single music collection. Another approach to discovering unknown music is simply adding all songs to a playlist and activating the shuffle mode. This more or less random playing of all songs can lead to serendipitous discoveries [51] and unpleasant songs can simply be skipped.

Cunningham et al. investigated the discovery of new music with 41 participants in a three-day diary study [39]. While they uncovered no unexpected sources for discovery, they were able to determine when discovery occured (surprisingly throughout all waking hours), where (mostly in residences or en route) and how (mostly via internet (21.8%) followed by radio (18.8%) and TV (13.2%)). More interestingly, they found that most reactions towards the discoveries were positive (62.3%) and happened passively (63.6%), i.e. without active involvement of the participants.

Overall, the most important source for discovering new music seem to be friends and acquaintances. A series of interviews by Laplante [94] showed that while online sources were popular amongst all interviewees, personal recommendations outweighed other sources for 10 of the 15 participants. Laplante found that "the quality of the recommendations ... seemed to be the main reason for relying on acquaintances rather than on formal music information sources" [ibid. p.115]. In both knowing and being known by the other person, recommendations are more tailored and estimating their value is easier. Sease et al. [146] found that the chance to receive social recommendations might depend on the size of one's own collection: "Participants with large collections rarely mentioned receiving CD mixes and regularly reported that they had few friends they could trust to introduce them to new music ... [they] commonly 'researched' new media through external sources or strangers - reading magazines, listening to the radio, or pivoting on 'like' items on Amazon, IMDB, or eMusic" [ibid]. Lee and Downey confirm the social aspect in their study with 427 people [97]: 84.6% ask friends or family for music information, compared to the next popular category 'record store staff' with only 45.7%. They also asked for places that triggered a music information search, and found again 'Acquaintance's or friend's place' to be the top result with 87.45%, followed by 'Public places' with 70%. Brown et al. [23] emphasize the importance of such "social music listening environments": "[They] promoted the exchange of information and taste about new music. Not only did friends get to listen to each other's music collections, but friends filtered music for each other, deciding what they thought others would like to listen to." [ibid.] One last aspect is that recommendations by friends also happen online via directed messages or more open publishing formats such as Twitter. This development is beginning to blur the boundaries between personal hints and blogging.

2.2.2 Acquisition

Having discovered a new and interesting song or album, people tend to include it into their personal music collection. Music is a special medium in that a single item is usually rather short and repeated consumption is desired (compared to other, more storytelling-based media such as books or movies) - at least if this repetition does not happen too often [40]. A personal collection gives listeners the freedom to access all of these items at any time.

Traditionally, music acquisition was buying a physical artifact such as a vinyl or compact disk and a music collection was also a collection of these artifacts. With the shift to digital music, music collections developed a more ethereal quality which is fitting with the recent discussions about network-based streaming via "the cloud".

At the moment, however, the digital transition has not been completed and there are still large amounts of physical music artifacts sold. BPI, the association of the British Recording Industry, announced in 2010 [21] that while the singles market in the UK was completely dominated by digital services (98% of all singles sold were digital), digital albums made up 19.6% of all album sales. Digital sales are trending upwards, however, so physical media for music appear to be a thing of the past. Another impending change is with the concept of ownership for music. While people still buy music "a la carte"[21] instead of completely relying on streaming it from the web, recent studies (e.g., [114]) show that this might change in the future.

Shifting to a digital format also led to easier copying of music and, in turn, an increase in illegal distribution. Music piracy, as it is dubbed by the music industry, is no new phenomenon, however. Brown et al.'s study from 2000 [23] that happened right after the changes introduced by the MP3 format and online file sharing, showed the pervasiveness of music sharing even with traditional media: "The mean amount of copied material in our enthusiasts' collections was 28%". The social aspect was again important as "the major source for material to copy came from friends" [ibid]. Also, cassette tapes were a popular way to create copies, even from CDs, or create all new mixtapes. The topic of piracy is controversially discussed by academia, industry and policy-makers and supporting studies can be found for either side of the argument. Oberholzer-Gee and Strumpf could not find a detrimental effect of piracy on music sales in a large-scale analysis in 2007 [122], while Rob and Waldfogel found a reduction on per capita expenditure in a sample of college students [139]. Peitz and Waelbroeck were able to explain the decline of CD sales in 2001 with piracy but not the one in 2002 [129] and Bhattachar-

jee et al. discovered an impact of piracy on chart ranking for everything but the most popular albums [17]. Unsurprisingly, the British Music Industry claims 1,200,000,000 illegal downloads in the UK in 2010 [21] and converts it to a nominal loss of £984,000,000. Despite the disputes about illegal downloading, no side denies that it is commonly used for music acquisition. 10 of Laplante's 15 participants, for example, admitted to "doing it frequently" [94].
Apart from illegal copying, multiple artists also make their music available online for free to promote themselves, a strategy which is especially popular with independant artists. Music blogs such as Pitchfork[3], aggregators such as The Hype Machine[4] or portals such as SoundCloud[5] or Jamendo[6] help with discovering new music and directly make it available. Also, some albums are "sold" in a pay-what-you-want model: Most famously Radiohead offered one their albums in this way in 2007 and studies by Kim et al. [81] show that this approach can succeed as a business model.
Finally, the latest development of online streaming services makes music acquisition obsolete: With a subscription comes direct streaming access to all licensed music of the service and with enough licenses a practically direct access to the "jukebox in the sky" (see figure 2.1).

2.2.3 Organization

Even though streaming services might make explicit music "collections" unnecessary in the future, having a clear idea of a definite set of one's own music will always be relevant and correspondingly tasks that surround it. One of Brown et al.'s interviewees mentions identity construction as a reason for having a collection: "'Even if you can borrow books from the library, it's still nice to have books. Because your library expresses who you are'"[22]. Having an organized collection is also a requirement for quickly retrieving music where suitable query terms are not known or available or for finding a complete set of songs that all belong to a certain category (e.g., a collection organized by musical genre can be used to quickly find a set of songs all belonging to the same genre).
Cunningham et al. performed an ethnographic study with ten of their students about music organization and report the results [41]: While most of their insights are based on music on physical media, they can just as well be applied to digital songs. One finding was the separation of collections based on the frequency of access: "Collections are generally divided into the active items (that is, those that see regular or occasional use) and the archival items ("music that is seldom or never listened to" [41]. This separation can also emerge spontaneously just through spatial proximity: "[Frequently used] music is typically located close to the playing device, to make it as easy as possible to quickly

[3] http://www.pitchfork.com

[4] http://hypem.com

[5] http://soundcloud.com

[6] http://www.jamendo.com

select the CDs for playing. The last played CD is usually placed on top, so that the less played CDs drift to the bottom of the stack" [ibid].
Bentley et al. also describe these two sets of songs: "'favorite' music (or music that is currently in favor) is kept convenient to the player" and "[some] music, like photos, gets 'archived,' and both end up in dusty collections that are rarely revisited"[16].
These findings might be an artifact of the physical nature of the participants' music collections. They also suffered from the problem of "geographic distribution"[41]: "In the car, for example, nearly all our participants had a tape player with a built-in radio. This limited their music listening to either the radio, or pre-recorded tapes. In the house, the standard setup was to have a main sound system in the living room, with satellite systems of limited capability in other parts of the house. This meant that some rooms, such as the kitchen or bathroom, became exclusively radio or tape playing rooms" [22].
Also in Cunningham et al.'s study: "The participants generally viewed having subsets of their collection in more than one spot as an annoyance, generally minor but occasionally major, since locating a desired CD might involve looking in many, sometimes widely separated, places"[41]. However, this problem and others related to physical media should disappear once a music collection has become completely digital. With wireless access to the internet, even storage restrictions of portable players disappear.
The organization style and the underlying logic appears to be mostly dependent on the owner: "Music organization styles and photo organization styles are also similar: some people organize meticulously, and others just throw things wherever they can fit them or wherever is convenient at the time... Labels vary from precise (date, time, event, and people for photos; descriptive title and track list for music) to cryptic (a photo titled 'the kids' included pictures of nephews as well as the participant's own children; a CD titled 'MP3 mix 1')." [16].
One of the rare studies focussing on digital music organization was performed by Vignoli. His in-depth interviews uncovered the preference for hierarchical folder structures in digital collections:"Four participants out of seven use a hierarchical structure based on artist-name/album-name/song-name with small variations, two based their classification on genre-subgenre and only one uses popularity for his collection... Almost all subjects use additional folders to store music that does not fit the structure of the collection. One participant uses language as the highest level of the classification... The way they organized their CD collection influences the organization of their digital music collections, basically they keep the same structure" [167]. Cunningham et al. also describe several different organization styles that they found in their ethnographic study: "by date of purchase, for example with the newest CDs placed either at the top of a stack or at the end of a shelf, by release or recording date, by artist, with the artists arranged alphabetically, by genre, where the number of genres can be large or small ('rap and other'), by country of origin (e.g., 'New Zealand music'), from most favorite to least favorite, in order of recency in which the CDs have been played" [41]. They also mention that "[a] secondary organization may be applied to each of the broad top-level categories (for example, sorting by artist within genre)" [ibid]. As mentioned above, music can also be organized based on intended use (for programming, driving, etc. (cf. [ibid.])).

Non-hierarchical systems such as labels can be helpful to get over the natural problems of hierarchies, such as items that belong to multiple categories (e.g., folk music with female vocals) or locating an item by browsing the category tree (cf. [150]). The beauty of organizing music based on folders and also the reason for its proliferation is that folder structures are based on the file system and can be easily transferred to other machines or portable players. In contrast, even though labels can be added to MP3 metadata within the respective section of the ID3 format[7], not all software or hardware players necessarily work with it which might make listeners reluctant to invest all the work of manually applying labels. Yet, organizing within a hierarchy is still more sophisticated than "simply adding in CDs to the top of a tower as they were purchased"[41]. Whether a collection is organized also depends on the involvement and enthusiasm of a person with music. Greasley et al. showed that music enthusiasts had ways of music organization beyond common approaches: "They made categorisations on the basis of their use of the music rather than its genre ('down there is the stuff you could handle in any mood' M:24yrs); on the basis of historical context or time periods ('stuff from the seventies' M:29yrs); on the basis of their own preference ('the less good stuff...the créme-de-la-créme' M:24yrs); or in some entirely different way ('I've got a different section for film music' F:24yrs)" [60].
The logic behind a music organization can also stem from identity construction: Cunningham et al. found "one participant [who] organized his CDs in a set of racks so as to allow him to hide some of his music: 'I can rotate my rack in a way that 'shows off' my best CDs while partially obscuring the average and embarrassing CDs'" [41]. Nick Hornby describes one especially eccentric example in his novel 'High Fidelity': "Tonight, though, I fancy something different, so I try to remember the order I bought [the records] in: that way I hope to write my own autobiography, without having to do anything like pick up a pen. I pull the records off the shelves, put them in piles all over the sitting room floor, look for Revolver, and go on from there; and when I've finished, I'm flushed with a sense of self, because this, after all, is who I am" [71]. This aspect of identity construction is especially relevant in collections with several owners: Sease et al. present a study on media collections in shared households [146]. While they found that platonic roommates' collections were not merged, sometimes even romantic partners refrained from doing it, due to reasons of identity and ownership. One participant explained: "I'm not sure I ever want them mixed just because it's... I don't know like it's... that's one thing that's all mine. We share a lot of other things plus the other stuff ... but that is just all my own" [ibid].
One caveat regarding the generalizability of research on music organization is the tempo of progress: Studies throughout the 2000s show the change from mostly analog collections [22, 16], to mixtures between analog and digital [41], to mostly digital [167]. Another influence on these results was the choice of participants: depending on whether participants were technology enthusiasts (e.g., from a lab or university [167]), they were also early adopters of new technology. The most recent study mentioned by Sease et al.

[7] http://www.id3.org

shows the general growth: Their criterium for a 'large' collection were 5,000 to 10,000 items (with some participants having large digital as well as physical collections) [146], while in Vignoli's study a large collection meant having only more than 1,000 digital songs [167].

2.2.4 Understanding

A related task to organizing a music collection is understanding it and learning about its contents. Underlying every music collection is a complex graph of connections between songs, lyrics and people. The bridge of a song might contain a subtle reference to another, a melody might be taken from a classical composer, and hip hop songs routinely consist of various samples, up to the extreme example of the artist 'Girl Talk' who gained fame with including more than 300 samples into his album 'Feed the Animals'[8]. Additionally, social networks of musicians, composers and producers are also reflected within song and album structures and finally all sorts of background stories related to bands and music.

As a first step towards understanding a music collection, having correct metadata is important. The default categories of artist, title, album, year, etc. are usually filled out correctly in digital music files directly bought from digital vendors such as iTunes. For music that has been 'ripped' from CDs, databases such as freedb[9] provide correct metadata based on audio snippets. Plug-ins for media players or external tools (e.g., TuneUp[10]) can also fix metadata of existing audio files.

For more advanced information such as lyrics or musical content such as key, tempo or pitch more advanced analysis tools (e.g., the echonest music analysis webservice[11]) or databases (e.g., musicbrainz[12] that even contains social relationships between artists) are required. This additional metadata presents some problems: it might be difficult to acquire (e.g., automatically "hearing" lyrics from the pure audio signal) and not every listener wants the same metadata: "The precise metadata desired is highly likely to vary from user to user, and so as rich a set as possible should be available, with the user able to select the fields of interest for display"[41]. This is also reflected in the ways people look for music: Bainbridge et al. analyzed around 500 music queries in an online reference system [8] and found that while 80% of all queries contained one of the basic information of artist or title, extended metadata such as genres, lyrics or circumstances of hearing were all equally likely to appear. Consequently, an ideal music information retrieval system would incorporate as many (and as accurate) metadata as possible.

Yet, some listeners might even want to go beyond these more song-related categories

[8] A visual breakdown of all samples on the album can be found on http://mashupbreakdown.com/feed-the-animals

[9] http://www.freedb.org

[10] http://www.tuneupmedia.com

[11] http://the.echonest.com

[12] http://www.musicbrainz.org

and understand their collection in-depth: "For [the friend of a participant], music does not begin and end with listening to the E.P., but continues onto a complete artist experience, including investigating the band on the Internet, downloading music videos and investigating their belief and social systems through thorough investigation of their official and unofficial websites"[41]. These music enthusiasts also scatter around dedicated online music communities such as last.fm, add information to wikipedia articles, and have been around since the early days of the WWW (Wanda Bryant's analysis of a folk music internet discussion group in her Ph.D. thesis from 1995 [24] is one example).
As mentioned above music collections have nowadays reached enormous sizes, sometimes even on mobile devices. Therefore, understanding all relationships and getting an overall picture of a collection requires an appropriate presentation. A variety of overarching visualizations for complete media collections have been proposed in the literature. A music collection with its metadata represents a high-dimensional data set, which makes it suitable for dimension reduction approaches such as self-organizing maps [85]. A common approach is therefore to display all items on an automatically generated map that reduces all available metadata (content- or context-based) to two dimensions. An early example were Pampalk et al.'s 'Islands of music' [126, 124] that generate a simulated terrain with islands representing dedicated sub collections. A subsequent version by Knees et al. extended the islands to the third dimension and added interaction capabilities [83] and Leitich et al. mapped it to a three-dimensional globe [99]. Other projects also rely on self-organizing maps but pass on the metaphor: The artist map [164] focuses on easy playlist creation and Stober and Nürnberger let the user adapt the influence of a dimension on the resulting map [157] or simply by dragging and dropping songs [156]. Self-organizing maps proved to be so popular within the music information retrieval community that a number of visualizations were also created to work in non-desktop environments, such as tabletop displays (e.g., [155], [69], [77]) or on mobile devices (e.g., [163], [113], [57]).
The common problem with these automatic maps of music is the non-descriptiveness of axes. A result of the dimension reduction is the merging of several dimensions with different weights on one axis, which means that the local neighborhood of a song contains "similar" items, but how this similarity is defined is unclear. Also, finding a song is only possible through direct search and not via navigating the map (except in a completely random fashion). Therefore, some systems rely on more comprehensible approaches that are easier to grasp for the listener, yet provide more powerful categorization schemes than regular consumer media players such as iTunes or the Windows Media Player. Torrens et al. apply a tree map visualization to a music collection - with similar problems regarding the accessibility - but also present a circular visualization technique, that uses circle sections for genres and the distance to the center as a mapping of the release date of a song [160]. Dachselt et al. present a so-called "facet-based browser" with *Mambo* [43], a zoomable user interface whose representation depends on the available screen real-estate or zoom level. All metadata are organized hierarchically and less important facets are only shown if possible. Another map-based visualization that allows more finely grained control over the characteristics of the map is *MusicBox*

[106]. Here, listeners can organize the songs based on various types of metadata, search and filter the resulting map and create playlists by drawing lines.
In general, however, most visualizations for music collections are still in a research stage and not in use by the general public. It is, of course, also a question of personal taste how much time is spent understanding and looking into one's collection.

2.2.5 Sharing

A defining factor within the interaction between people and music is sharing it with friends or acquaintances. As a part of the social functions of music (cf. [64]), sharing in its most basic sense only means making someone aware of something. Sharing can happen due to various reasons: for supporting a musician by enlarging their audience, for wanting another person to enjoy good music, to establish oneself as a part of a group or emphasizing one's musical identity, and also for creating a self-image of knowledgeability and showing off. Whatever the reasons, sharing is prevalent with both physical and digital representations of music.
The invention of the music cassette and the subsequent home-taping movement was the first instance where copying and sharing of music entered the mainstream consciousness [23]. Copying music became easy and sharing good music with friends commonplace as it was also a natural extension of music discovery: "Friends also often searched through each other's collections, looking for music that they might borrow or perhaps copy" [ibid]. A special case of sharing are shared collections: "While individuals have their personal music collections, they may also participate in a shared collection with others for example, students sharing accommodations may keep a stack of CDs by the living room stereo, or families may have developed a shared collection that everyone can contribute to and play" [41]. Sease et al. had an in-depth look at such shared collections and found that they are usually very open: "If limits were established with family members, it involved age appropriate material for children." [146]. Participants were, however, aware of the identity impressions they made with their collections, but "... on the whole were willing to accept that judgments were made" [ibid].
Finally, one last version of sharing physical media is the mix- or compilation tape culture mentioned above, that forms a cross-over between simple sharing of existing items and craftsmanship: "In general, the role of friends in recommending and acting as 'guides' for new music had an important influence on what people later bought or listened to. One way this happened was through the swapping of compilation tapes made up by friends (see also (Willis, 1990)). This way of sharing music, although time consuming and cumbersome with most technology, was particularly valued" [22]. Nick Hornby also lets his narrator in 'High Fidelity' describe the artisanship required in creating a good mix tape for wooing prospective partners: "A good compilation tape, like breaking up, is hard to do. You've got to kick off with a corker, to hold the attention (I started with 'Got to Get You off My Mind,' but then realized that she might not get any further than track one, side one if I delivered what she wanted straightaway, so I buried it in

the middle of side two), and then you've got to up it a notch, or cool it a notch, and you can't have white music and black music together, unless the white music sounds like black music, and you can't have two tracks by the same artist side by side, unless you've done the whole thing in pairs, and . . . oh, there are loads of rules" ([71], p.45). Cunningham et al. explored these rules by analyzing requests by playlist enthusiasts from the website *Art of the Mix*[13] [38]. They found the importance of having a central theme around which the playlist revolves. This theme can be based on a genre or artist, on romantic songs (similar to Hornby's narrator), for accompanying an activity, sending a message or also as a challenge to the playlist builder: "a mix of songs with 'eye' in the title" [ibid].
Mix tapes are also used for story-telling and reminiscing and strengthening social bonds: "As with photos, our users often made special music mixes in order to say something special to someone or convey parts of their own life to others. Chris made a CD for his brother with songs of importance throughout his brother's life for his birthday"[16].
Little surprising, sharing culture and reasons have not changed much with digital music, but technology has on the whole made it easier to copy and transfer music. While using a tape to copy music could take a timespan up to the original duration of the source, digital copies are created in an instant. This accelerated not only illegal copying (see above) but also made spontaneous sharing much easier and even allowed sharing music with the whole world - within limits: Sease et al. discuss this aspect for their study: "All participants were comfortable sharing their media and their media identity with those with whom they are more intimate. Sharing with friends, colleagues and absolute strangers happened less frequently, but if participants felt that they had special knowledge or access to certain media, they were more inclined to share with the world" [146]. Another reason for publicly sharing music is its rarity: Some of Sease et al.'s study participants describe a feeling of duty to make rare recordings available and find the resulting copyright infringement acceptable, as the rightsholders appeared to ignore their own duty for doing it. "The more unique the item, the more a collector feels a need to share outside their immediate realm of intimacy, illustrating an ethical dilemma between open access and fair use versus ownership rights" [ibid]. One major difference between online and offline music sharing is the option to stay anonymous. While this certainly restricts sharing for personal identity management and is, as some argue (e.g., [90]), one of the reasons for the success of (illegal) file-sharing, it still allows pushing an artist's fame by making rare recordings and bootlegs available or overriding regional restrictions enforced by rightsholder associations or record companies. Sometimes, even artists themselves circumvent restrictions by their labels: The band *OK Go*, for example, encouraged fans to watch their latest music video on a different online video platform than Youtube, as their label *EMI* had prohibited the embedding functionality for private websites there - contrary to their own wishes[14]. Closer to encouraging copyright infringement, *Danger Mouse* instead of a record sold a book with photographs inspired

[13] http://www.artofthemix.org

[14] http://okgo.forumsunlimited.com/index.php?showtopic=4169

by his album *Dark Night of the Soul* and a blank CD-R with the instruction to "use it as you will" after legal disputes with his label[15].
Non-anonymously, the digital world also has equivalents to shared personal collections as can be seen in Voida et al.'s study on iTunes music sharing in an office network [169]. Similar to the shared physical collections in Sease et al.'s study [146], sharing in this case was no active process and relied more strongly on the concepts of impression management. Other interesting findings were that some participants also felt that having a certain musical expertise required them to invest work into their collections and curate them, similar to the owners of rare items. Voida et al., however, also emphasize how iTunes' interface and other technical aspects such as subnets within the network impact the sharing behavior.
With digital music, the place of mix tapes has been taken over by pre-defined playlists that can be imported in media player software (e.g., iTunes) and listened to, provided that all songs are available in the collection. This problem of availability is less prominent for streaming services and playlists created for them are accessible and ranked by the community on various websites (e.g., ShareMyPlaylists[16] for the Spotify service).
Spontaneous sharing of single songs has become much more convenient with digital technology: Hyperlinks to songs can be sent via email or instant messaging or made available for dedicated circles by posting on social networks or one's blog. Websites for musicians' self-presentation such as myspace[17], bandcamp[18] or Soundcloud[19] provide tools for the usual sharing via social networks, but also for embedding player widgets in personal websites, so visitors can instantly listen to a song. Dedicated platforms for music enthusiasts like last.fm allow demonstrating one's knowledge about music by displaying a suitable listening profile (see above and [153]).

2.2.6 Listening

The final and most important step and the actual reason for all curating actions around music collections is listening. First, the right song has to be retrieved. The chances of succeeding in finding it are dependent on the degree of organization of a physical or digital collection and also on its "geographic distribution"[41] and other technological limits (cf. [22] with mostly physical music). In the digital realm, however, music information retrieval systems (for an overview see [46]) can help in finding a song based on singing or humming a part of its melody (*Query-By-Humming*) or other songs that are similar to a given item either in their musical qualities (such as key, mode or tempo) or other aspects (such as artists involved, lyrics or the context in which they were last

[15] https://www.eff.org/deeplinks/2009/05/danger-mouse-inducem

[16] http://www.sharemyplaylists.com

[17] http://www.myspace.com

[18] http://www.bandcamp.com

[19] http://www.soundcloud.com

listened to).
Most of the time, however, a person is not interested in retrieving song after song to listen to them, but instead opts for defining or using a pre-defined playlist or filling the play queue of their media player. In Cunningham et al.'s study about *Art of the Mix* requests [38], motivations for building personal playlists were "as background for another activity, ... to reflect a particular mood or emotion in the creator, ... [or] as 'party music'" [ibid]. Predefined playlists are either music albums where the order has been determined by the artists themselves or movie soundtracks and compilations that contain music by various artists. Listening to complete albums provides a less laborious way to listen to music than actively selecting single songs for play queues or playlists. Oftentimes, however, even such granular control over the music that is played is not required, so other strategies can replace dedicated selection: Leong et al. [100] argue that shuffle and skipping make up for the shortcomings of player software regarding extremely large music libraries and the resulting detrimental effects on the user experience: "Shuffle is a good example of an approach that is suited to supporting interactions whereby users' goals are fuzzy and further goal definition would require too much effort" [ibid]. Using shuffle can also play into discovery by leveraging serendipity, as mentioned above. An important difference in this shuffle listening compared to other random processes is that listeners can restrict the amount of randomness and the frequency of required skipping by applying shuffle only to a subset of their whole music collection. In a different study, Leong et al. examine online self-reporting on music listening [101]. They found that in 91 of their 113 cases listeners used shuffle (in 69 of these even on their whole collections), while in 15 of the remaining sequential cases an album was used.
This observation is in line with Bentley et al.'s discussion of the similarities of photos and music that showed that people usually exercise a behavior that relies on the notion that a musical choice just has to be "good enough": "Participants often just played what was already in the player. Gina: '[I'll] check what's in here first. If there's anything cool in [the player].' ... In the same way, on the computer users would just play the current play list or they would pick music that was lying out already" [16]. The authors connect this listening strategy to the ideas of *browsing*, *satisficing* and *sidetracking*. Browsing describes a search without a clear goal (when listeners are unsure what they want), while satisficing is picking the first item that fits roughly and sidetracking triggering a change in goals caused by a finding during the satisficing process. Sidetracking also happened during listening to music: "For instance, if a song by Weezer came up, they might decide that they wanted to play the entire album" [ibid]. Brown et al. also observed the differences in listening behavior: "Some enthusiasts felt strongly about the pleasure of looking through the spines and selecting CDs to play. Others preferred to avoid this, and often left CDs in the player for days on end" [22].
Low-involvement interactions such as skipping can also be used for measuring the appeal of a song. Pampalk et al. describe building a recommender system based on skipping, by removing songs with a similar content [125]. Skipping behavior, however, happens in bursts, similar to other human behavior [9]. Bentley et al. observed "... that

although our participants might let many songs play before skipping one, once they skipped one song they often continued skipping several more times before they let a song play again. Our explanation for this behavior is that they lowered their aspiration level in order to avoid having to go to the effort of skipping a song. Once they had started skipping it was easy to keep doing so, and they raised their aspiration to a more discriminating level" [16].
The most consequent way of spurning music choice is using either music programs designed by other people or automatic systems. Radio transmissions were used for listening to music even before the proliferation of personal music existed. Due to the internet, they are nowadays no longer restricted to their transmission range but can be received practically everywhere, which gives listeners a detailed choice of predefined programs. Broadcast radio is, however, not necessarily a listener's first choice. Brown et al. observed in their study: "It was clear that for some participants the radio served as a back-up source of music when they ran out of pre-recorded music since they would often be in situations where they only had access to a limited number of tapes (such as the car) or CDs (such as the kitchen)" [22].
One common approach for receiving better playlists without investing work are personalized systems based on recommender systems [2]. Recommender systems are peculiar information retrieval tools insofar as their goal is not clearly defined: While regular IR tools can claim to retrieve right or wrong items, the results of recommender systems are on a scale between more and less suitable. Still, they fit very well to the task of retrieving music, as a "right" song only has to be "good enough" [16]. Music discovery by recommender systems happens in several ways: First, on e-commerce websites such as Amazon.com where listeners can be inspired by suggestions for their favorite items [94] ("Customers Who Bought This Item Also Bought"). Second, websites such as last.fm create listener-independent similarities between artists based on collaborative filtering of all of their collected listening histories. Third, playlist generators such as Apple's Genius are also based on recommender systems and suggest other suitable songs for buying. Finally, the most prominent use of recommender systems online are personalized webradios such as Pandora[20]: Based on overarching artist or tag radios they gradually adapt to the listener's preferences through simple interactions such as skipping songs [125]. They also provide advantages for discovering new music as no active interaction is required (in the optimal case where the system only selects suitable songs) and that songs are directly heard instead of visually displayed.
Comparing radio stations and recommender systems, however, shows that people still trust humans more than machines: In a study by Paul Lamere (reported in [48]) people were asked to rate playlists that had either been created by a radio DJ, an algorithm or randomly and also guess how the playlist came about. The results showed a high tendency of the participants to give good ratings to playlists that they thought had been created by a human (on average 3.33 of 5, compared to 2.76 (algorithm) and 2.08 (ran-

[20] http://www.pandora.com

dom) respectively), even though the actual winner was the algorithm with a rating of 2.63 (compared to 2.64 for random and 2.49 for the DJ).

2.3 Music listening studies

In the previous sections I have shown the various reasons for listening to music and the tasks that accompany music collections. In the following, I will give a survey of studies on music listening. As shown above, there are several reasons that make generalization in this regard difficult:
The first is the rapid change in technology and thus human behavior surrounding music. Results from a ten year old study might provide only small insights into today's listening reality and developments that are happening at the moment, most prominently the change from personal collections to digital streaming services, will also radically change human interaction with music.
Second, both people and music are very different individually and studies can only shine a light on certain people and certain music. All results are necessarily anecdotal, even though technology might allow us to take much larger samples of music listening in the future (see the listening history study in the next chapter).
Finally, there exists an abundance of different ways to listen to music and listening to it is not necessarily a deliberate decision. Sometimes, circumstances might require listening to it even though it is against the listener's intention and taste.
In this section, I will focus on *active listening*: The listener is a self-determined actor in this scenario and has free choice which music to pick (within the limits of her or his collection) and when to listen to it. This excludes several scenarios of music listening that were discussed in section 2.2.6 such as radio or recommender systems.
This section gives an overview of types of studies and their methodologies for exploring this topic. These studies provide input for the description of listening factors in the next section.

2.3.1 Design space of listening studies

All studies presented here stem either from musicology, sociology, cultural studies, computer science or human-computer interaction. Despite their diversity, all these disciplines have an interest in understanding human music use. Acquiring information on listening is always based on observing or asking actual music listeners, but the situation in which this observation is made can vastly influence the insights gained.
Studies about music use can be categorized according to five different factors (see also figure 2.2):

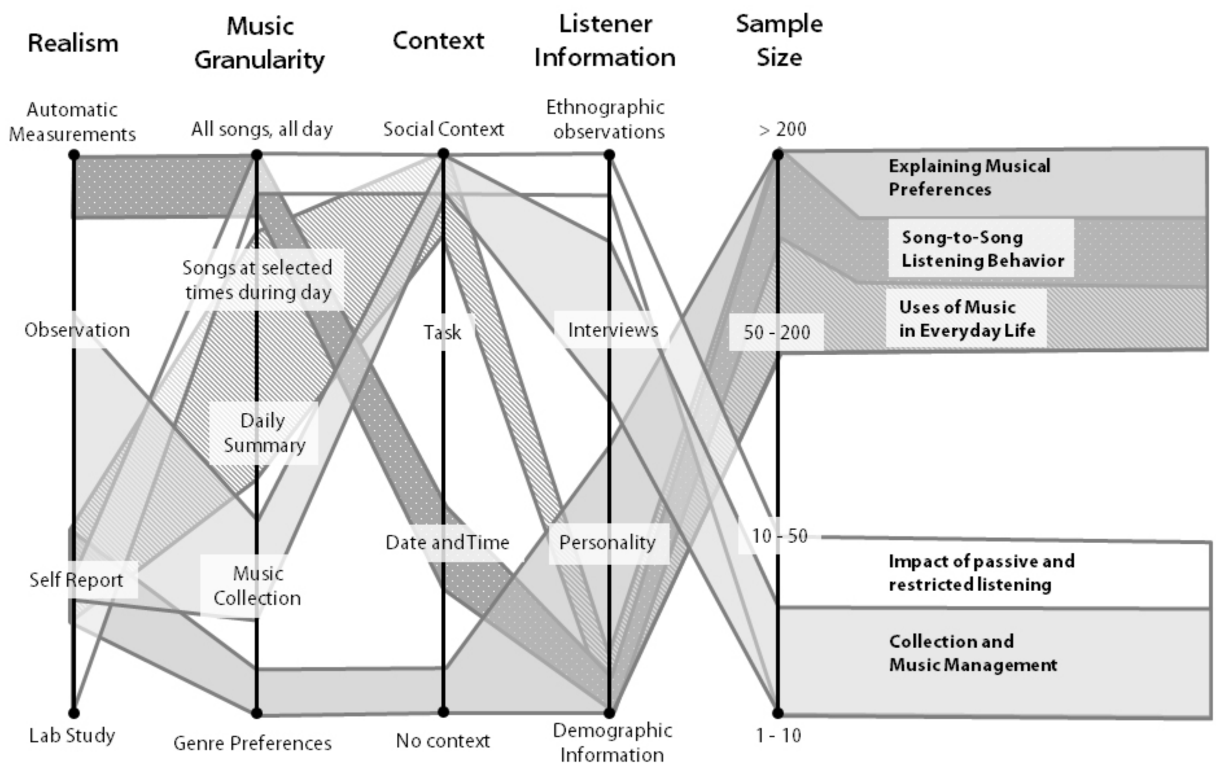

Figure 2.2 Scientific studies about listening behavior can be categorized along five dimensions: (1) Degree of realism, (2) Granularity of captured music, (3) Listening context, (4) Information about participants, and (5) Size of the sample.

- *Realism:* Real-world situations make it hard to control relevant variables in an experiment. Therefore, to measure the impact of music on moods and performance researchers often approach it in a laboratory study with full control over the environment. On the other hand, if we want to learn more about the uses of music in everyday situations, studies have to restrict a participant's behavior in as little ways as possible. Another aspect is if music experience is reported by the participants themselves or selected/measured by the experimenter. In the former case, social desirability bias or simply forgetting can skew the results.

- *Music granularity:* Not only when and how somebody listens to music, but also what they are listening to is very interesting for researchers. It might, however, be difficult to capture the minute-to-minute interaction with music, especially if study participants are asked to document their music listening themselves. Therefore, some studies simply ask for general preferences regarding genres or favorite artists without delving into the everyday listening behavior.

- *Context:* A central concept of music listening is context, as the current listening context might determine what task is to be supported with music and if music is appropriate at all. Therefore, context can be the only explanation for music choices, even if they are captured on a very detailed level. A simple version of context is the current date and time of day that already allows guessing if the participant was working, at home, etc. More complex descriptions take a listener's actual task or even the social context into account.

- *Listener information:* Another aspect that can be described in various degrees of detail is the listener her- or himself. Almost all studies rely on demographic information such as age, gender and social status, as these have proven to be vastly influential to music behavior. Some studies go beyond that and also integrate a listener's personality with corresonding personality questionnaire such as the Big Five inventory (see e.g., [136]) which allows looking for correlations between character traits and musical taste. The most costly studies try to learn as much about their participant's lives and motivations as possible, mostly in the form of in-depth interviews and long-term ethnographic approaches.
- *Sample size:* The larger the amount of information required from participants, the more laborious it is to collect it. Therefore, studies who go in-depth and explore participants' lives and music habits also often make this effort only for a small number of people. The downside of small samples is that they might not be very representative or overemphasize statistical outliers. A small sample size can also be caused by requirements that are too large for regular participants (e.g., documenting every song heard in a day). Technical support that automatically logs parts or all of the relevant data can help with increasing the possible sample size.

These dimensions can be useful to understand the meaning and expressiveness of a study. The following studies will be categorized according to these aspects.

2.3.2 Listening studies

All researchers planning a study about music listening have an idea on what they want to find and what aspect they want to explore. As music listening is so prevalent, a single study can only uncover so much, even if it is only about active music selection. Depending on the sought-after results and the scientific background of the researchers, the methodology also changes.

Impact of passive and restricted listening: A first group of studies looks at the impact that (mostly passive) music listening has on people. While these cases are not about active choices, they give more insight into the actual effect of music on emotions and mood. For a completely controlled experience and thus indisputable results a lab studies are usually the best choice. The classical studies by Konečni et al. [87, 86] conform to these criteria and control everything from participants' moods to musical complexity. Similarly, studies measuring the actual processes of the brain by using, for example, MRI technology (e.g., [108, 61]) are also highly controlled. While this control leads to high accuracy and intricate insight into musical attributes and human response, they also reduce the realism of the results. As musical stimuli are simple and the environment is artificial, one might argue that these experiments not necessarily capture the broad appeal of music. Also, the time-consuming preparations and limited access to required technology keeps sample sizes small.

More realistic approaches of measuring music's impact are mostly concerned with and played out in commercial situations [116, 115, 121, 177]. The methodology is relatively straightforward: Form a hypothesis on musical impact, play or let participants choose respective music, let shoppers/dining guests/yoga students shop/dine/exercise and measure the sales numbers, training results or which music was chosen. Due to the constant stream of (unknowing) participants, these types of studies can reach a high number for sample size, but are again restricted in their expressiveness regarding real listening situations as experimenters mostly choose the music or prepare the selection. Still, these studies are highly relevant for enterprises, as they show how music can be used to influence decisions (even though there is no 'silver bullet' music that increases sales).

Finally, studies that explore the use of music in daily life are oftentimes based on regular sampling throughout the day and also take passive listening episodes into account (e.g., [120, 44, 154]). Similarly to the controlled studies above, participants do not have control over the music that is chosen but are still subjects to its impact which might have a more negative tendency. In all other regards, these studies are identical to the 'Uses of music' studies which are explained below.

Collection and management studies: A second type of studies is concerned with the collection-aspect of music and tasks surrounding it. These studies are often performed by library researchers or computer scientists, as they are most interested in the ways people interact with large collections. Tasks within collections are similar to the ones presented above (see section 2.2) and include acquisition, ordering and retrieval.

To learn about ways people interact with music, most commonly an ethnographic approach in combination with interviews is chosen (e.g., [41, 169, 16, 146]), as human reasoning and motivations are too complex to capture with, for example, a simple questionnaire. Also, learning the organizational principle of a collection by simply looking at it (for example in a shared iTunes collection [169]), would most certainly fail due to ambiguity, partial disorder and the influence of the owner's background. In these interviews and corresponding home and collection tours, people can explain themselves and their music and what logic, ideas and stories they applied to their collection.

Collection studies usually discuss the aspect of listening to music, but only in the context of music organization (for example, keeping piles of the most popular CDs near the stereo [41, 146]). They provide deep insights into people's motivation and feelings about music, but require much time and are therefore mostly restricted to tens rather than hundreds of participants. In terms of listening behavior, they also only draw a very abstract picture and their results might be distorted due to desirability bias.

Explaining musical preferences: Musical taste is different across people and sociologists and music psychologists aim to find out what aspects dictate this and what patterns might be hidden. These studies mostly aim at measuring broadly defined characteristics of people (demographic or personality traits) and mapping them to taste in music. Due to the large variety in these aspects and musical taste, such studies take large-scale approaches to find any patterns. They usually rely on questionnaires as study instruments and categorize music listening according to genres (e.g., [138]) or self-defined overarching categories (e.g., four dimensions in [136] or five in [135]).
A popular line of research started by Rentfrow and Gosling in 2003 [136] explores the connections between personality and music traits (examples are [31, 181, 32, 135]). The approach is relying on an inventory for assessing personality (usually the Big Five Inventory [75]) and one for taste in music and derives connections between them. One notable study that overcomes the problem of bias in self-reports is by Dunn et al. [47], who additionally use three months of automatically tracked listening behavior: While self-reports correlated with listening behavior, they also had the general problem of these studies which is the fuzziness of the genre concept (cf., [7]).
Other popular research targets with large samples and questionnaires are gender (e.g., [89]), social status (e.g., [62, 118]) and age (e.g., [142, 140, 145, 20]).
As these studies are focussed on uncovering underlying large effects, their results allow guessing (rough) musical taste from demographic or other easily collectable information. The description of music is abstract and details about the actual listening activities, however, scarce (listening sequence or changing of taste over time, for example, are not taken into account).

Uses of music in everyday life: The line of research that is concerned with the uses of music in everyday situations is also very close to active listening and capturing realism. The most basic approach simply asks people for their music taste and presumes it as static and correct: Questionnaire studies about socio-demographic factors and music (see above) take this approach but can only make statements about high-level observations.
A more realistical way to measure the impact of music on people in their everyday lives is to observe and talk to them. Tia DeNora took such an in-depth approach in her book 'Music in Everyday Life' [44] and tells fascinating stories about people's lives with music. The ethnographic approach, however, also requires much effort which restricts the sample size. Brown et al. describe in a purely interview-based example [22] the impact of the then new MP3-technology on music listening. Yet, interview studies can only tell so much about music as they present only a singular and overarching version of music, mostly aggregating the context. Also, they reflect only the most impressive situations, as more mundane ones are likely forgotten shortly aftwards.
To capture the true everyday music, listeners can be asked to record what songs they currently hear throughout the day. As this is laborious to do manually and might even keep participants from listening to music (to avoid having to write it down), samples are usually taken only once or a few times at random points of a day (cf. Experience Sampling Method [37]). North et al.'s study from 2004 [120] and Greasley et al.'s from 2011

[59] used text messages once or five times a day, while related studies notified pagers [159, 154] or palmtops [78]. These studies give an intricate picture of the role of music in people's lives, overcome problems of memory and even describe the impact of passive music listening on their participants. They can also focus on specific concepts such as favorite songs [93] or explore personality and music preferences (young people's use of rock music [159]) using this methodology. As participants themselves do most of the work, sample sizes with several hundred participants are possible. This burden on participants, however, is also a downside of the approach as it might cause evasion or skewed results and they are also plagued by the usual biases of self-reports. Taking only a few samples a day is also not detailed enough to talk about the role of sequences, playlists and shuffle in music listening.

Song-to-song listening behavior: A last type of studies that has only recently taken hold investigates the strategies people employ in listening to music and looks at music on a song-to-song level. Capturing this detailed listening behavior would be too time-consuming and error-prone to do manually and would almost reverse observation and activity (with songs lasting a few minutes, documenting listening would take almost as long as actually doing it). Therefore, these studies rely on automatic capture which requires suitable technology and is not perfect (see section 3.3), but produces much more accurate results than self-documentation. The main downside, however, is that besides date and time more complex descriptions of context are lost. These methods are therefore not yet fit for capturing uses of music in everyday life, but can be used to explore massive, context-less samples on a very detailed scale. As most studies that employ this method are also closer to science than the humanities this lack of context is not problematic for the results and analyses are commonly purely quantitative.
Performing such analyses also does not necessarily require the consent of participants and researchers can simply use public listening data from the internet. Such "big data" analyses [4] can uncover interesting relationships: Lambiotte and Anderloos extract social groups and genres from listening histories [91], Zheleva et al. derive a model incorporating listening sessions into taste descriptions and recommendations [178], and Park and Kahng derive temporal listening patterns of a community of over 70,000 people [127]. For our own take on a large-scale listening analysis see section 3.4.
A slightly different approach with similar goals can be found in Leong et al.'s studies. They looked at song sequence and especially shuffle listening, but used conventional approaches such as an analysis of various (informal) online posts [101] and a one-week long diary study with 12 participants [100] as input. This also gave them the context and motivations of their participants, while suffering from similar restrictions as mentioned above.
One general downside of studies on detailed listening is that they can either only describe a whole community or need a highly involved listener (cf. [60]): If individual listening histories are just created by pressing the shuffle button then their expressiveness is more about the randomizing algorithm than taste.

2.4 Listening factors

As we have seen, music has different functions in everyday life, requires certain tasks due to the transition from a performance to a collection medium and different aspects of music listening are studied in different ways. Music listening is, however, such a complex and varied topic that creating an overarching, all-describing model is difficult if not impossible. Still, the large amount of research on it allows at least restricting the influential factors.
In this section, I will present an overview of relevant factors that control active listening behavior based on the literature reviewed above. These factors answer two main questions:
1. When, i.e., in what situations, do people listen to music?
2. What music are they choosing?
There have of course been other models of listening before: Rösing describes [142] a scale based on the attention of the listener (from low to high cognitive involvement). Ross based his reception model on the *person*, the *product* and the *situation* ([143], cited after [142]). This is similar to Hargreaves and North who postulated that: "In the broadest sense, an individual's response to a given stimulus or piece depends on an interaction between the characteristics of the person (such as age, gender, musical training, personality); of the music (e.g., its complexity, familiarity, style, etc.); and of the situation in which it is encountered"[64]. Leong et al., however, concentrate on the listening process and distinguish between *listening mode*, *content organisation*, *activities during listening* and *affective outcomes* [101]. These models, however, ignore the technological aspect and overarching patterns from global trends.
In the following overview, all listening factors are organized into three categories that represent different levels of context: *Personal context* includes factors that are mostly relevant to the listener her- or himself and might not be openly accessible. *Local context* describes social and situational aspects of the immediate surroundings of the listener. The *Global context* finally incorporates the state of the world as a whole and represents overarching influences that are relevant for large parts of the listening population.
All of these factors influence listening decisions but they fall into two categories: Some have a *gradual* influence with varying particularity and determine the type of music that is to be heard in different levels of detail. Others, however, only regulate the access to music and are strictly *binary*: They completely restrict or completely allow access to a certain song.

2.4.1 Personal context

Factors of the personal context represent aspects that are only changing very slowly, if at all. They represent the backdrop to music listening and their gradual versions only give very vague predictions. The impact of personality, for example, might be determining what genre of music will be listened to next, but not which artist or specific song.

Demographics

Demographic factors such as age, gender, country of residence or social class influence music decisions on a more abstract scale. They might not allow predicting specific decisions, but give a general idea how a person's taste in music might be shaped.
Especially in a young *age*, musical taste changes regularly and radically (cf. [65]) before settling for a more general and open listening style (cf. [146]) or just listening to popular music (the prevalence of so-called 'Adult Contemporary' music that is strikingly uniform on German radio stations is a sign of that).
While there is no clear effect between *gender* and genre choice (e.g., [47], except for heavier music such as heavy metal and rock in adolescents [65]), Kreutz et al. argue [89] that women have a more empathizing than systemizing listening style than men but this is independent of the chosen genre.
The *cultural background* provided by one's country of origin also influences music listening: Musical choices are informed by the native language (with English as music's world language as a fallback), the importance and acceptance of music, and also the quality of music education (Boal-Palheiros et al., for example, discuss the role of music in British and Portuguese schools [20]).
Social status has a heavy influence on other factors such as education and access to musical training which in turn determine music listening and a tendency towards high- or low-brow cultural engagement [62]. In a large-scale study with 2,500 participants, North et al. found the impact of these other demographic factors such as relationship status, employment and travel preferences on music listening [117, 118, 119].

Personal taste, preferences & music involvement

Another group of personal contextual factors describes personality, musical preferences and the importance of music in one's life. Again, these factors will only gradually and, in the case of music knowledge, effortfully change and can make more accurate predictions about which and if music is used.
As described in sections 2.1.2 and 2.3.2 above, *personality descriptions* also are reliable predictors for music taste: Extroverted people like to listen to energetic and rhythmic music, shun conservatism, and so on [136]. Together with the various other predictions of personality for political leanings and other areas of life, personality can already give a solid overview of a person.
A second factor are *musical preferences* that heavily influence music choice. Even if the personality makes rough predictions about preferred music, preferences shaped by a lifetime of experiences and stories provide the strongest motivation for choosing certain music or not. This also includes irrational prejudices (disliking a certain artist or a certain genre on principle: "This is girly music!") and self-perception (see section 2.1.3

about identity construction as a central use of music).
The *intellectual capabilities* of a listener also determine her or his taste in music. This follows from the 'mood management' function of music (see section 2.1.1) that attributes different values for arousal moderation depending on the capabilities of a listener: Quirky pop music just might not cut it for a fan of sophisticated jazz and an opera by Schönberg might be overwhelming for the uninitiated listener. Therefore, cognitive attributes (that also only rarely change) are a good predictor for tendency towards high- or low-brow music.
Similarly, if a person had *musical training* can have an impact on taste. Playing an instrument requires not only learning and understanding musical concepts (and thus the intricacies of more advanced musical pieces), but also listening to and spending a lot of time with music. Unsurprisingly, children who had musical training more often state a liking for 'serious' music [65].
Finally, the most important criterium for the intensity of music interaction is *personal involvement*. If a person is interested in music or not determines whether they take the time to learn about new releases, read the backstories of their favorite artists, dig through record collections, find rare gems, and most importantly: listen to it. A person with a high involvement can be a dedicated playlist tinkerer who even creates suitable transitions for a mix, while a low involvement listener would pick one of his five CDs, press Shuffle and be done with it. As time spent with doing something increases knowledge about it and the tendency to continue doing it in the future (see *effort justification* in social psychology), people with a high involvement for music will probably never lose it, and can be predicted to be more knowledgeable, eclectic in their choice, and overall listen to more music at more opportunities [60]. Therefore, simply asking someone if they're into music can already provide a large deal of information about their listening behavior.

Access to music

One last set of personal context factors concerns the actual access to music. These factors are mainly binary and closely linked to the technology factors (see below). At least negative predictions are made easy by knowing about the access factors: A song that is not available cannot be listened to.
A first aspect of these factors is the *size and variety* of one's music collection. This mostly depends on the enthusiasm towards music and the effort brought up to find items for and maintain a large collection. But it is also determined by socio-economic factors, as large collections are either expensive or have to be collected in garage sales and thrift shops (see e.g., [146]) which takes time. Simply downloading massive amounts of music from the internet might fill the collection, but does not solve the problem of organization (thus rendering some parts of it obscure or inaccessible - the "shoebox effect" [16]) and also poses the risk of a law suit.
Technological advancements such as internet video portals like Youtube, that are increasingly used for music consumption (e.g., [21]), or *music streaming services* like Spotify

can play a role in easing this access to music by providing millions of songs at one fell swoop. One has to be aware, however, that these services only provide access, mostly through text-based queries. Without knowing which music is good, finding it is practically impossible and has to be supported using recommender systems, social networks or traditional music-centric media.
Access to music also increasingly depends on the capability of a listener's *mobile device*. Since the Walkman, music players have limited storage capabilities and can only carry so much music (even if an iPod Classic with 160GB capacity might hold a lot of songs). And while this access restriction through limited storage might be eased in the future with always-on mobile data access to one's favorite streaming service, other restrictions remain on the go: Technical failures, areas with no reception or simply an empty battery can quickly limit one's access to music from millions of songs to none at all.

2.4.2 Local context

The factors in the local context are highly instable but much more relevant for predicting music consumption than the constant personal ones. While the personal factors determine the overarching, general taste in and tendency towards music that only changes very slowly, the local context has a much higher fluctuation: It changes in the worst case from song to song (see musical context below), but mostly stays stable for at least one listening session (the temporal context can last much longer, however). In return, the local context also provides a much higher accuracy in determining musical outcome.

Temporal context

Even though people might not be aware of them, our lives are filled with temporal cycles and patterns whose effects mostly disappear in the noise of everyday life. Listening strategies often also adapt throughout days and weeks and fulfill the mood management function of music. The longer temporal cycles become, the less predictive they get. So seasonal cycles might be good for predicting the appearance of Christmas carols, but say less about general styles throughout a certain day.
The most prominent of temporal cycles is the *sleep/wake* cycle that has also a prominent effect on music listening. Most people switch off their listening devices when they go to bed (even though some like to listen to audio books for falling asleep).
The *time of day* also determines if music is heard and what type of it. Some people like to start their days with cheerful and energetic music and end it with more mellow tunes, in order to manage their arousal. Also, certain playlists are helpful for supporting tasks such as work (see below) that often occur throughout the day, while in the evening more entertaining or party music might be played.
A type of rhythm with a lower frequency distinguishes between *weekdays and weekends*. As several studies (e.g., [120]) have shown, people listen to more music during the work week (often for supporting their tasks or while commuting) and spend their weekends

away from music playback devices.
Finally, the longest types of cycles look at the impact of *seasons* on music listening, where annual holidays such as christmas and the corresponding speciality songs appear. Also, summertime might bring gaps in listening similar to weekends, as people go on vacation (even though music listening might merge forever the songs with the impressions they got during the vacation). The season can also determine the general type of music that is played: Some people match their music listening to the seasons, playing relaxed music during the summer and darker songs during winter.

Emotional context

Another local impact on listening behavior comes from the current emotional state of a listener. This again is a reflection of mood management through music. Emotional states tend to change slowly throughout the day and depending on the situation and the source they can linger for longer or shorter. This is also an effect of personality traits (more cheerful persons are not as heavily and as long sad as more gloomy ones).
A first use for music in the mood management case is the *reversal* of negative moods or the increase of arousal through energetic music. Some people seem to have an instinctive feeling for what music can help them in managing their arousal and play according one.
There is also a certain overlap between the personality traits and this mood management aspect. The tendency towards reversal is not universal and depending on their personalities, people tend to either *reinforce* the mood of the current situation through music (depressing music for a sad situation) or counter its negative effects with it (uplifting music for a sad situation).
Once the general tendency of a listener in this regard is known, the emotional context can be used to predict music listening depending on the mood of the song. Interestingly, there does not seem to be a connection between "chronic emotional states" and musical preference, as Rentfrow & Gosling have shown [136].

Musical context

The local aspects of the musical context capture not only the relationship between a person and a song, but also cover the songs that came immediately before in the listening session (which is not applicable for the first song in a session, of course).
Listeners share personal histories with their songs that might be different for everyone. While one person heard a certain song during a special time in their lives and remembers it every time they hear it, another might find the same song annoying due to its cheery pop nature. These memories and stories make music a perfect tool for *reminiscing* and also have an effect on musical choice. But apart from personal connections with events, song relationships also develop over time, so after a phase of playing a song repeatedly, fatigue might set it and the listener avoids it until a happy rediscovery much later [41]. Depending on where in this *lifecycle* of the relationship the song currently is, a person

might be more or less inclined to put it on.
The musical context also describes what songs came immediately before the last one. This *immediate history* of songs is probably the most solid tool for predicting songs that will immediately follow: Having listened to the first five tracks of an album, the chances are great that the listener will continue with the sixth one (similar with other types of pre-defined playlists). Also, it is usually a safe bet that the listener will not listen again to the song that just played, except when that song already had been played two times in a row (as the listener seems to have grown fond of it). Apart from music choice, once a person has skipped one song, chances are that she will continue to do so (cf. [16]). For one, all of this behavior regarding the immediate history can be derived from the general burstiness of human behavior [9] - people have a certain inertia in their actions. These predictors also usually reflect the rules for good playlists (see section 2.2.5).

Task

A major influence on what and if music is played is the current task. Similar to the mood management aspects explained above, music can also be used to make a boring task more interesting or take stress of an intense activity.
So, a first factor influencing music choice is simply the wish to fight *boredom*. A large number of tasks such as housework or driving a car do not require full focus and are therefore underwhelming and not challenging enough for most people (cf. Csikszentmihalyi's theory of flow [37]). They therefore support them with music to help ease the monotony. In North et al.'s study on the uses of music in everyday life [120], "'It helped to pass the time' was more likely to be a reason for listening to music during the working day than during the evening".
On the other hand, music can also be a more *active support* for doing certain tasks. Exercising is regularly named as an activity that people support with suitable music, for example rhythmic and energetic songs in a certain tempo range. Such songs are also used by external actors to increase productivity or the enjoyment of people (e.g., in an aerobics class). To return to North et al.'s study [120], in 21.6% of all cases when participants were on their own, music helped them to concentrate/think and thus supported their current task. It is also characteristic that people create playlists and name them after certain tasks (cf. [167, 41]).
Music listening can, however, also be inversely affected by a certain task. In some contexts it might not be *appropriate* to listen to music, either not to disturb other people (even though this reason is becoming increasingly obsolete due to headphones) or for reasons of appearances (a store clerk should be able to spontaneously interact with customers without first getting rid of the music gear).
Finally, the task can also be used not only to predict which playlist is chosen but also to describe the probable *interaction* with music. While exercising, people usually like to concentrate on their own body instead of fumbling with their music players. Therefore, certain tasks lend themselves to more complex interaction (adding songs to the current

playlist while commuting on the subway) while others are too intense to be suitable (demanding traffic situation).

Social context

A last set of local factors describes the current social context of a listener. Similar to the current task, if and which people one is with makes a strong case against or for certain music (or no music at all).
Music is commonly consumed while *alone* and without other people (this can also mean an anonymous situation as common in big city life). In such situations, people are generally unrestricted and able to freely decide if they listen to music and which songs they choose. In lieu of social interaction, music then also commonly becomes instrumental in fighting boredom with the current task.
When listening in a group, the current situation and the social relationships determine the role of music. As is the case with the current task, some situations deem music *inappropriate*. This is usually the case in offices during meetings or other activities with more than one person, where listening to and interacting with one another is mostly hindered by additional auditive signals. Also, if music is allowed, choice over it might be very restricted: During a funeral service, for example, solemn music is important while cheerful music is (usually) taboo.
In more relaxed social situations such as parties, music is an essential component. In such casual scenarios, there is often no explicit DJ but every guest is free to bring their own music choice into the selection which makes it a *shared playing*. Usually, people can either make use of the host's music collection or even connect their own music players to the stereo. Everybody is free to change the music, but songs should be played to the end and the host usually reserves the right to override a selection or end some guest's role as a DJ. The to and fro between actively selecting and passively listening also influences one's own music choices: Some song that a party guest plays might trigger an association and lead to the wish of listening to some specific other song.
Finally, it is important if the listener is playing music *for somebody else*. This might be for the party crowd in the above scenario, where especially for adolescents choosing certain music can determine one's standing in the group, or in a more intimate context to impress a date or create the right atmosphere. In a fringe case, the listener might also be paid to work as a DJ and has to entertain a bar or a club. Here, musical choice can also be determined by explicit wishes from the crowd.

2.4.3 Global context

Listening factors in the global context are overarching aspects that are only very rarely under the control of a listener but influence her or him nonetheless. These represent the impact of society, mostly in form of cultural and social habits and fashions and technological developments. As habits change only slowly and mostly in response to technology, global factors are, similar to personal ones, not very specific in their predictions.

Commercial influence

Commercial influence has an impact on access to certain music and determines whether a listener might be aware of its existence. On a more fundamental level, however, record labels act as *gatekeepers* for music. They determine whether a new band gets a contract, the financial support to develop their music and record an album. Their promises can give young musicians trust in their own talent and the financial outcome or get them to quit music. While this strict role of major record labels has degraded in the last years with the success of smaller and independent labels and the various ways for artists to promote themselves and sell their albums on the internet, not every musician wants to or is able to handle the business and public relations side in parallel. Without intrinsic conviction of their own ability, artists still like to rely on the experience and the financial support of a label.

A second major factor that is often ignored in research is the major influence of *presentation* on the shaping of musical taste. As described above, music has an important function in identity construction and songs and artists stand for certain messages and ideas. These are, however, usually not only communicated through the style of music and the lyrics, but also through the appearance of the band members in promo shots and music videos, interviews and additional media such as websites. Impressive stage shows with eccentric costumes or gory photos in a CD booklet can make artists more or less attractive to certain groups of the population. This image is also the reason why some listeners might enjoy songs by Lady Gaga or Justin Bieber but never admit to it and file it under "guilty pleasure" (they would also probably delete it from their listening histories, cf. [153]).

Commercial or political interests are also the reason for explicit *restrictions* in access to music. Labels might decide to release certain albums not or only with a certain delay in certain regions of the globe. Streaming services receive licenses for songs only region-wise, which leads to the inconvenient situation that British Spotify users might be able to listen to and share certain songs with their Swedish friends, but they cannot listen to them on the same service. Performing rights and other societies that represent artists might restrict access to songs due to failed licensing negotiations, leading to the dreaded "This video is not available in your country" warning on Youtube. Artists themselves might also want to restrict access to their music: It took almost ten years until the remaining Beatles allowed digital sales of their collected works on the iTunes music store. Finally, governments can also censor or restrict access of minors to "inappropriate" songs and videos, where the inappropriateness is usually debatable. The effect of all these restrictions is often forcing fans to either abandon the artist or resort to piracy.

Trends and culture

Overarching music trends are determined by a society at large. Especially regarding identity construction, culture says which music is for whom which is enforced through peer pressure. Such *genre trends* are even relevant for people who are actively shunning this social pressure, as there are certain genres more suitable for rebellion than others. Also, genres that are labelled as "mainstream" create a negative, but still existing, effect on people who avoid trends. These trends also work on a larger, commercial scale. As a record label's primary reason for existence is to generate shareholder value, labels also follow cultural trends and support them. Each music genre has a target audience and the repertoire of a label should have something for every customer. Due to commercial interests, niche genres can also enter center stage. The most prominent example of this is the history of rap music, that started as a medium of expression for social outcasts and developed in less than thirty years into one of the most commercially viable musical genres.
Influential for a single music listener is also the *role of music* in a society. This is determined by the general treatment and respect for music and also the implicit mapping between certain genres and demographic factors. Music can be seen as either a form of mindless distraction or valuable artifacts of culture. Also, the stereotypes inherent in music listening lead to preferences or ignorance towards certain genres: German Schlager is only appropriate for a certain age category and younger people normally only admit to liking it to emphasize a conservative or traditional mindset. Similarly, there exists music for each combination of gender, sexual orientation, social status and age group and knowing about the delicate differences between music is important in order not to convey a wrong image.
A last influence of society on listening comes from the definition of *appropriateness*. It is convention that allows and sometimes even demands music in certain situations while it is rude in others and these rules also commonly change. An example for that is the rise of portable music players that started with the Walkman as tolerated nuisances inflicted by teenagers and by now have become commonplace for almost all age groups in public transport and other locations. Convention also not only determines if music can be played but also what type of it. Listening to Christmas carols for the holidays instead of Surf Rock is simply a rule that has been transferred from generation to generation and has nothing to do with the songs themselves.

Technology

One last and very important set of listening factors describes technological aspects of listening. The massive change that technology has brought in recent years to music consumption could be seen in the above examples and also in the datedness of some of the research.
The *novel ways of listening* through all forms of new devices and services make music a much more central part of our lives than it was before. This in turn also influences the role of music, makes it a commodity but also affordable and accessible for more peo-

ple. As described above, novel technology such as online video platforms, digital music stores and streaming services have radically changed access to music and portable devices have made music available everywhere.
Technology also allowed *changing the interacting with music*. Online platforms give music buffs a platform to discuss about their favorite artists, social networks allow immediate sharing of one's new music discovery and platforms like last.fm enrich one's online profile with listening histories. Recommender systems and recommendation-based web radio such as Pandora merges listening with discovery. And creative music apps for mobile devices provide new ways to access one's music collection and create playlists.
A last technological factor are *monthly costs* that come in the form of fees for streaming services. Paying for access can influence behavior, as music suddenly becomes cheaper (relatively) the more you listen to it. Formerly one-off payments for CDs turn into month-by-month renting. Also, mobile music listening via streaming can be influenced by the mobile data plans of a provider: Bandwidth limits can keep listeners from accessing music either to conserve it for more important tasks or because the volume has already been used up towards the end of the month.

2.5 Summary

In this chapter I have shown the influence of music in our daily lives and the state of research on it. At first, I described based on research from music psychology and cultural studies the three main categories of uses for music, which can be summarized under the terms mood management, cognitive functions and identity management. I then gave an overview of tasks surrounding music which included discovery, acquisition, organizing and understanding a collection, sharing and listening to it. These tasks were mostly found by researchers from computer science and other technology-oriented disciplines as they are to construct the tools for supporting music consumption. Afterwards, studies about active music listening, their scope and methodology were presented. In conclusion, I derived all relevant listening factors that influence the when and what of music selection from the literature.

Chapter 3

Music Listening Histories

Sell, I'll sell your memories,
For 15 pounds per year
But you can keep the bad days.

– Muse - *The Small Print* –

When Charles Babbage designed the Analytical Engine in the 19th century, he also added capacity for keeping track of up to 10,000 numbers. This earliest modern computer had its store mainly for keeping track of calculations or before print-out. The introduction of various sensors, the inclusion of other non-mechanical forms of storage and the increase in available space brought the computer's version of memory much closer to its biological counterpart and allowed a dream of extending human capacity for remembering with external tools: Vannevar Bush's visionary 'memex' [25], a microfilm-based device similarly shaped to a desk, was supposed to store not only books and magazines and replicate associations through an early version of hyperlinks, but also personal conversations and thus keeping track of one's own activities.
As all technological advancements seem to progress logarithmically it took another sixty years until Bush's vision finally came to life: In the *MyLifeBits* project [15], Microsoft researchers Gordon Bell, Jim Gemmell and Roger Lueder tried to capture a lifetime's worth of media and communication and provide convenient ways to access it (fortunately, they did not have to rely on mechanical apparatuses for retrieval). But they went even further than that. Bell, who acted as the team's guinea pig, increased the available information by also wearing a SenseCam at all waking hours, a device that took snapshots at interesting moments, and capturing conversations using a tape recorder. During the project, the researchers kept up with the development of storage and bandwidth and gradually increased the fidelity of the electronic memory. In this regard, MyLifeBits

went much further than the original Memex idea: Bell and Gemmell had nothing less in mind than "a complete digital record of your life, a complete e-memory of your time on earth... by recording your life digitally you have the opportunity to bequeath your own ideas, deeds, and personality to a posterity in a way never before possible. With such a body of information it will be possible to generate a virtual you even after you are dead" ([15], p. 6). In this vision of perfect *lifelogging*, every aspect that is captured by our biological senses and stored in our biological memories will receive a perfect digital replica in a way that is not affected by forgetting.

The appeal of this vision that promises to make it easier to recall names, events and conversations and finding overarching, subtle patterns and not to forget its benefits for personal narcissism made Bell et al. pioneers of the lifelogging movement. As portable devices and sensors also obey Moore's law and thus become more powerful and cheaper every year and hold more and more storage, an army of hobbyists and practitioners is by now exploring the implications of tracking one's every move. Especially measuring health aspects and improving nutrition and exercise through sensors and logs is a popular past-time and only recently the movement has postulated their goals in the first official *Quantified Self* conference (the community is also still unsure about the actual terminology as *quantified self*, *lifelogging*, *life tracking*, *personal informatics*, and other terms all seem to describe similar and overlapping concepts).

Lifelogging, however, is much more prevalent than one would think and has already reached the mainstream (without them being aware of it). Services such as social networks, personal blogs, instant messengers, search engines, photo sharing sites, etc. all capture little parts of our lives and even today contain already broad pictures of ourselves. And as media consumption also constitutes an important aspect of one's life, there are also services for tracking books, movies and songs, producing detailed reading, watching and listening histories.

In this chapter I will first give an overview of research on lifelogging, on sensors for capturing the data, and technology and interfaces for accessing it. Afterwards, I describe the characteristics of an ideal listening history based on the important listening aspects that were determined in the last chapter. I will then present the current state of listening tracking and discuss the restrictions and downsides of the current approach. To show the impact of collecting large numbers of listening histories on understanding listening behavior itself, we performed a large-scale analysis of such data and present the results and insights. The chapter closes with a mapping between the ideal listening history and its real-world counterpart and suggestions on how to fix the differences.

3.1 Related Work on Lifelogging

The vision of lifelogging encompasses everything that a person can experience in their lives, so the problems of capturing and working with such data are remarkably similar over all different categories. With music listening histories being a specific type of lifel-

ogging data, they also benefit from understanding the larger lifelogging problems and perspectives.
Li et al. propose a model for lifelogging data [105], that composes of five stages: *Preparation*, *Collection*, *Integration*, *Reflection*, and *Action*. While geared towards health-based systems (especially regarding the *Action* phase) and rather data-centric, this model gives an overview of the different phases a lifelogging system should support. In a similarly overarching way, Sellen and Whittaker describe the possible benefits lifelogging can have [149], something which is not as obvious as it seems. Existing research mostly focusses on challenges around capturing and preparing data, while the benefits are simply taken for granted. Also, systems for accessing lifelog data lack an overarching understanding of what exact use case they are supporting and which ones they are missing. Sellen and Whittaker call their classification *Five Rs*: *Recollecting*, *Reminiscing*, *Retrieving*, *Reflecting*, and *Remembering Intentions*. Recollecting and reminiscing both describe working through sections of the past, while retrieving goes for a specific piece of information. Reflecting (equal to Li et al.'s definition of it) makes overarching themes available, and the remembering intentions use case allows for planning and other future activities.
One interesting aspect of lifelogging research is that most of the progress is not driven by academia anymore but appears in a more ad-hoc fashion. Online sharing of device designs and experiences, the low-cost foundation of suitable online services, sensor-laden smartphones, and startups that produce affordable devices have led to a large community of non-researchers experimenting, mostly on themselves. Work on the different parts of the lifelog cycle can be subsumed under the technological perspective of collecting and storing accurate representations of the data and making use of and accessing it.

3.1.1 Collecting and storing lifelogging data

How lifelog information is collected mainly depends on its type: The three main cases are capturing of purely *digital activitives*, of physical activities with all-round devices such as *smartphones*, and capturing of physical activities with *dedicated devices*.
Capturing *digital activities* is most of the time trivial. Each digital device already captures interaction and other processes and using this data for lifelogging is simply a question of storing and accessing it. Web browsing histories, email- and instant messaging data, and interaction histories with media such as photos, videos or music are usually directly available for working and playing with them. Last.fm's *Audioscrobbler* software (see below) is an example for this approach that accesses common media player software for tracking music listening. Another development that increases the ease of capturing digital activities is the transition from dedicated desktop-centric computing to online- and cloud-services, as a new form of mainframe computing. With multiple devices per person and the subsequent need for synchronization, more and more activities happen online anyway (Google's Chromebooks - a line of always-online laptops without a local filesystem - are the current end of this development). Most web services provide

easy-to-use application programming interfaces (APIs) that allow third parties to develop suitable tools that make use of this data. Therefore, digital activities can usually be captured automatically without bothering their initiators.

While a lot of interesting activities happen in the digital realm, actual human experience mostly takes place in the physical reality. Lifelogging apparatuses have to take that into account, but require suitable sensors to capture it. As these sensors do not necessarily exist yet (e.g., perfect detection of social context) or it is difficult to extract high-level information from their low-level readings (e.g., the owner's movements from carried accelerometers), certain parts of lifelogs still have to be created manually, with all the downsides regarding accuracy and overhead. Still, technological development promises to reduce such necessities and even refine current low-level data (e.g., by detecting food or faces in existing photos and video streams).

A low-cost approach for either documenting or capturing lifelogs that relies on a device that many people already carry everyday are *smartphones*. They provide online access and multiple powerful types of sensors (gyroscopes, accelerometers, GPS, camera, microphone) and can easily programmed without additional costs for the lifelogger. A first set of smartphone apps falls into the category of self-reporting. Apps to support diets (e.g., Lose It[1], Calorific[2]), sports (e.g., DigFit[3], RunKeeper[4]) or general-purpose tasks (e.g., Daytum[5], TallyZoo[6]) are all available for popular smartphone platforms such as iPhone and Android. More sophisticated applications rely on on-board sensors for tracking sleep (e.g., Sleep Cycle[7]) with the integrated accelerometers. Other apps use these motion sensors for tracking steps and running activity (e.g., iTreadmill[8]) or GPS data for hiking (e.g., MyTracks[9]). Research projects also commonly rely on run-of-the-mill smartphones for creating lifelog prototypes. Hwang and Cho, for example, present a Bayesian network approach to automatically detecting landmarks in phone-based sensor streams [73].

Sometimes, certain sensors are either not available in regular smartphones, not accurate enough or would cost too much battery life to be useful. A last category of lifelog capture therefore works with *dedicated devices*. Before cheap multi-purpose sensors and devices were available, researchers had to build them themselves. Pioneering work in lifelogs and wearable computing came from Steve Mann, who explored the notions of sensor-laden clothing and wearable cameras (e.g., [107]). Similar to Microsoft Research's extensive SenseCam project [70] or the competing inSense project [19], pictures were taken

[1] http://loseit.com

[2] http://www.worksmartlabs.com/calorific/

[3] http://digfit.com

[4] http://runkeeper.com

[5] http://daytum.com

[6] http://tallyzoo.com

[7] http://mdlabs.se/sleepcycle/

[8] http://www.itreadmill.net

[9] http://mytracks.appspot.com/

in regular intervals or when sensor-readings suggest something interesting happening. A SenseCam, audio recorders and various manual approaches like scanning documents were also used by Gordon Bell and colleagues for his lifelog [53]. Other early approaches like Forget-Me-Not [92] relied on a ParcTab or other PDA-like devices for storing and accessing lifelog information. Of course, researchers also use combinations of existing sensor-technology and devices for experiments (e.g., Oliver and Flores-Mangas used a heart monitor and a smartphone for testing their exercise-based music-selection system [123]).
Commercial approaches all use the health angle for selling their products: Fitbit[10] or Zeo Personal Sleep Coach[11] either rely on motion or brain activity and track steps per day or sleep. Nike+[12] captures running activity and energy monitors such as the OWL[13] keep track of home energy consumption. With all the available commercial and home-brew offers, developments like the ANT+ wireless sensor standard[14] promise to stop the chaos of devices and hacks.

3.1.2 Accessing/Using lifelogging data

Once lifelog data has been collected the question often remains how to most effectively access it. Sellen and Whittaker's Five Rs [149] provide ideas what the benefit of this information might be, but making sense of it often requires dedicated software. Approaches for accessing lifelogging data fall into one of three categories: They are either *fully automatic* and happen without any involvement by (and sometimes even awareness of) the lifelogger, in a *structured* way, that categorizes data but provides no aggregation, or as in *summarizing* form with overviews or visualizations. Usually, data is available in a low-level format that allows all three approaches.
Fully automatic lifelogging access usually happens without the owner of the data being aware of it. Examples are recommender systems that shape music and media recommendations based on previously consumed items. Similarly, websites use cookies to capture web browsing histories and provide additional services (e.g., filtering search results or making suggestions). Also, news streams in social networking sites or news aggregators are shaped by the surfer's behavior and interests derived from her or his lingering and clicking on certain items.
Structured access relies on a suitable automatic tagging and retrieval of single items. Digital lifelog data such as photos taken with smartphones usually also carry time stamps and location information. More advanced sensor setups can also capture other information such as the wearer's posture or the audio surroundings (e.g., inSense [19]). Lifel-

[10] http://www.fitbit.com

[11] http://myzeo.com

[12] http://www.nikeplus.com

[13] http://www.theowl.com/

[14] http://www.thisisant.com

oggers themselves can also add tags manually afterwards or mark and add annotations for interesting situations (e.g., audio annotations with the Fitbit tracker). This wealth of annotations and structural information is necessary to aid in retrieving from and understanding the lifelog. Biological memory might be fallible but it is also notoriously flexible: almost any type of memory can be used for describing some other through associations. When trying to recover the content of a discussion, the colors of curtains and the smell of the flower bouquet might be valuable information. Lamming and Flynn were going for a prosthetic memory device with their Forget-Me-Not prototype [92] that could capture as much context as possible to make it easier to retrieve suitable information. Similarly, the MyLifeBits project also spawned a complex user interface and boasts "[support] for 25 item types. Each item has around 20 common attributes. For example, contacts have an additional 62 attributes, including email address and date of birth" [53]. The interface allows searching and filtering for single attributes or combinations thereof. Web-based services for collecting lifelogging data commonly rely on direct search and time-centric approaches for accessing information. One exemplary approach is Memolane[15], a service that aggregates data bits from various services and combines them into one timeline. The idea of organizing personal information along a timeline, however, stems from Freeman and Fertig's *Lifestreams* prototype, which pioneered it in 1995[52].
A last set of access tools uses *Summarizing* approaches. Summarizing aims at the *Reflection* use case and provides an overview of the data and insight into patterns that might not be obvious on the level of single items. Aggregating data points can happen on different levels of complexity: Health-based services usually provide a summarized version of calorie intake or average heart rate and other top-level statistics to motivate the owners of the data. Similarly, services might be counting the occurence of certain events or the general size of a lifelog data collection. Using visualization instead of simple numbers can also be helpful to grasp the content: The daytum[16] lifelogging service allows creating free-form *displays* that show certain aspects of the data or small summarizing charts.
Time lines are helpful in visualizing personal conversations and adding the importance of certain contacts can facilitate understanding social relationships [165, 166]. The *LifeLines* project [132] also relies on a time line but uses various biographical data (their examples contain medical and court records) to show overarching patterns and understanding histories.

3.2 Ideal Music Tracking

Music listening histories are a special case of general lifelogging data, which means that the possible benefits as well as the difficulties in collecting and accessing this data are

[15] http://www.memolane.com
[16] http://daytum.com

similar. The main goals of collecting listening histories can be subsumed under the three categories of access above, namely *fully automatic*, *structured* and *summarizing*. While automatic approaches rarely need more than the number of relevant items, the more metadata and contextual information is available, the better structured access works. As described above, varied attributes can be helpful in retrieving a piece of memory and especially with music, more metadata means being able to cater to the taste of more people (cf. section 2.2.4). Beyond these memory triggers, summarizing and understanding the data requires local context. The listening factors (see section 2.4) rely on various extra-musical information which means that simply collecting songs and time stamps is not enough. The more information is available and the easier it is to mentally go back to the point in time where a song was listened to, the clearer listening decisions become: Some of them might only make sense in the social context, others simply stemmed from an empty smartphone battery or a maxed-out data plan.
Another aspect that is relevant for music lifelogs is who is working with them after collecting them. If the creator of the history her- or himself is trying to make sense of this information, certain aspects can be stripped, as people tend to be aware of them. Factors from the personal context such as demographics are usually static enough to be left out for a personal history. Also, listeners' memories can be helpful in making sense and reflecting back on certain periods of their lives - something that is not available to an outside observer. So if somebody else is using the history to understand listening decisions (for an example see below, section 3.4) all listening factors are important.
Based on the listening factors and the goal to maximize available contextual information, data from an *ideal* music listening history can be separated along the lines of Personal, Local and Global context.

3.2.1 Personal Aspects

Personal context only rarely and then very slowly changes, so one description per person is usually sufficient. Relevant descriptions are:

- *Demographics:* Demographical information (age, gender, country of origin/residence, social class, etc).
- *Personality:* Descriptions of personality such as Big Five Inventory results.
- *Musical preferences:* Favorite genres or artists.
- *Intellectual capabilities:* As part of personality, results from an IQ test, for example.
- *Musical training:* How much and how intense musical training has the person had.
- *Personal involvement:* How much time is the person spending on music, are they enjoying it, etc.

- *Size and variety of the personal collection:* Similar to *personal involvement*, the number of songs available over time and the relationship to an "average" collection at that time.
- *Access to music streaming services:* Has the listener been online at a time. Has she or he generally had access (e.g., via a subscription) to such a service.

3.2.2 Local Aspects

A second set of aspects represents the local context of listening. As this changes rapidly (usually song-to-song or session-to-session) keeping track of it helps both creators and observers of a history. The current status should also be kept for each song. Local aspects encompass the following factors:

- *Date and time:* When was a song listened to. This information can also be useful in contrast with other parts of the history to see overarching patterns (sleep/wake cycles, time of day, weekdays & weekends, seasons).
- *Emotional context:* The current emotional state of the listener at the time. Preferably enriched with prevalent thoughts and ideas.
- *Musical context:* Information to make the song identifiable and as much metadata about it as available. Also, the listener's opinion of and associations with the song at that point in time and its current position in the song lifecycle. Finally, the source of the music (portable player, desktop software, streaming service).
- *Task:* The listener's immediate activities at the time of listening and her or his intentions (also in regard with the music). Also includes the description of the situation to gauge appropriateness of music or the amount of interaction possible.
- *Social context:* Who the listener was with at the time and was the music partially or fully directed towards somebody else (e.g., as a DJ).

3.2.3 Global Aspects

The last group of aspects contains global contextual factors. They also remain fairly static, similar to the personal factors, and have the advantage of being applicable to larger parts of the population. Available technology and conventions are usually similar for more than one listener. Global aspects depend on the following factors:

- *Commercial situation:* The impact of large and small record labels on music production and promotion. Relevant ways of advertising and its reach.

- *Trends:* What genres are in and out or mainstream and underground (depending on the personality of the listener). The role and acceptance of music in everyday situations.
- *Technology:* Costs and availability of listening devices and services. Access to novel tools for discovery and organization. Factors impacting mobile music access such as mobile data plans and their costs.

All these aspect constitute parts of an ideal listening history. Such a perfect history would contain all the information that is required to make its owner's behavior regarding music completely transparent for her- or himself and any outside analyst.

3.3 Real-world Music Tracking

As should have become clear during the description of the attributes of an ideal listening history, such a collection of data would reach far beyond the scale of a list of songs. Fortunately, such a level of detail is not necessary to already reap some of the benefits of a music lifelog and parts of the ideal data are already collected in commercial tools and services. This section describes the available providers, the structure of their respective listening histories, and the shortcomings and problems of the resulting data.

3.3.1 Music tracking services

Real-world services and tools that track listening behaviour fall into two broad categories: They either collect listening only *locally*, i.e., within their own confines (e.g., media player software that tracks songs that it played). Or they follow a *global* approach and try to capture all songs that the respective listener has heard.
Local collectors are media players in software or device form such as Apple's iTunes software and iPods devices. Webradios and -sites that allow listening to music also frequently capture and store this behavior locally. These features allow playing back the songs heard last (similar to the stack of last CDs next to the stereo [41]), seeing when a song was heard last and also the number of times it was played over the lifetime of the application. These are the main goals of this data collection, while the lifelogging aspect is usually ignored and no access to the data provided. Such lifelogs would also be seriously restricted as only listening from one source is contained.
On the other hand, *global collectors* aim for a complete music listening history that spans all possible sources. Such an approach faces challenges, as not only all available software media players have to be taken into account, but also proprietary hardware players, smartphones, and other sources of music. One solution to this diversity is manual

collection of the data. The media recommendation service GetGlue[17], for example, uses this approach for any media from movies, to books and music. While manual collection is suitable for media with a longer duration, actually keeping track of each song that is played is too tedious for most listeners, highlighting only the favorites and not the complete history.

One of the pioneers in the area of global automatic collection is last.fm[18]. Their service has been available since 2002 and mainly promises to act as a webradio and provide suitable recommendations. This, however, requires an overview of a listener's taste in music. Initially, last.fm only relied on manual feedback in the radio player, where listeners can skip songs (a subtle way of expressing dislike, cf. [125]) and 'love' them. Shortly after their launch they started integrating listening data from Audioscrobbler, a company based on a daemon software that tracked listening in various media players - an act they called *scrobbling* - and subsequently merged their companies. Due to the available API[19], tracking in other players did not necessarily have to be added by the Audioscrobbler team themselves and widespread support soon made last.fm the de-facto standard for listening histories. Also, the web-based nature of the API meant that no dedicated Audioscrobbler process was necessary and other applications could directly transmit listening information to the last.fm servers. Therefore, scrobbling is also available, for example, on non-computer devices (e.g., Logitech's Squeezebox music players), handhelds (Spotify's mobile client also supports scrobbling and even caches listening items when no network is available), and from web-based sources (ex.fm[20] provides an interface to online-MP3s and also updates the listening history). Enthusiasts also opened more exotic scrobbling opportunities: *Scrobbyl*[21] takes audio information from vinyl and other analog sources, identifies it using audio matching and sends the result to last.fm. The success of last.fm also spawned various competitors with similar approaches such as libre.fm[22] (more control over the resulting history), tunewiki[23] (location-based music trends), or like.fm[24] (better support for web-based music sources).

Still, last.fm boasts 30 million active users a month [25] which makes them the current market leader in this field.

[17] http://getglue.com

[18] http://last.fm

[19] Application Programmer Interface

[20] http://ex.fm

[21] http://scrobbyl.com/

[22] http://libre.fm

[23] http://www.tunewiki.com/

[24] http://like.fm

[25] R. Jones. Last.fm - The Blog: Last.fm Radio Announcement, http://blog.last.fm/2009/03/24/lastfm-radio-announcement

3.3.2 Raw listening histories and metadata

Naturally, the listening histories produced by these real-world services are not nearly as complex as the ideal version discussed above. However, they already contain enough information to provide useful applications. In the following, I will take the example of last.fm which represents the minimal version of a listening history containing for each listening instance only a *song identifier* and a *time stamp*.
The *song identifier* provides a unique description of the song that has been listened to. In last.fm's case this contains the artist's name and the name of the track. They acquire this information either through the provided ID3-metadata of a song or audio fingerprinting, where a short acoustic sample of the song is sent to last.fm to be compared to a database of existing songs [26]. Last.fm also tries to correct typos in a track or artist name and attaches terms such as '(live)' for live-versions of songs.
The *time stamp* of a listening instance contains the time and date a song has been listened to. Providing the correct time zone for the listening instance is not possible without locating the listeners, so this information is kept internally in the UTS (Unix Time Stamp) format, which provides the seconds since midnight January 1st, 1970 in the UTC$\pm$0 time zone. Each time stamp has thus an accuracy of one second and contains the starting point of the song instance.

One peculiarity of last.fm's listening histories is also that they do not contain the duration of the listening instance. This stems from the original intention of capturing listening histories as input for a recommender system that only requires positive or negative votes for each song. This intention also led to the Audioscrobbler's default behavior of only logging a listening instance if a song has been heard for at least 30 seconds (i.e., no skipping has occurred until then). This setting is however adjustable.
Having a uniquely identifiable song means that various other data sources can be accessed to recreate some of the contextual information that is necessary for benefiting from listening histories (for the following cf. [13]). One common scheme for this metadata is the hierarchy of songs, albums, artists, and genres (see figure 3.1).
Organizing something as diverse is music can be difficult but the sheer number of songs requires some form of classification. "A common (but vague) abstraction that combines content- and contextual information is the musical genre. Genres such as 'Rock', 'Pop' and 'R'n'B' describe not only a defining sound and style but also a certain context of the music (e.g., 'Brit-Pop'). Genres are hierarchical in nature ('Alternative Rock' is a sub-genre of the more general 'Rock') and this hierarchy is commonly extended to artists and their albums"[13]. Genres are, however, not without their problems [7]. They are not necessarily accurate or clearly defined (a problem for studies based on self-documentation by the participants, cf. [135]). Applying a strict hierarchy of genres, artists, and albums also means that one artist necessarily belongs to one genre which restricts such a classification to a musically-stable set of artists. Yet, implementing a

[26] R. Jones. Last.fm - The Blog: Audio Fingerprinting for clean Metadata, http://blog.last.fm/2007/08/29/audio-fingerprinting-for-clean-metadata

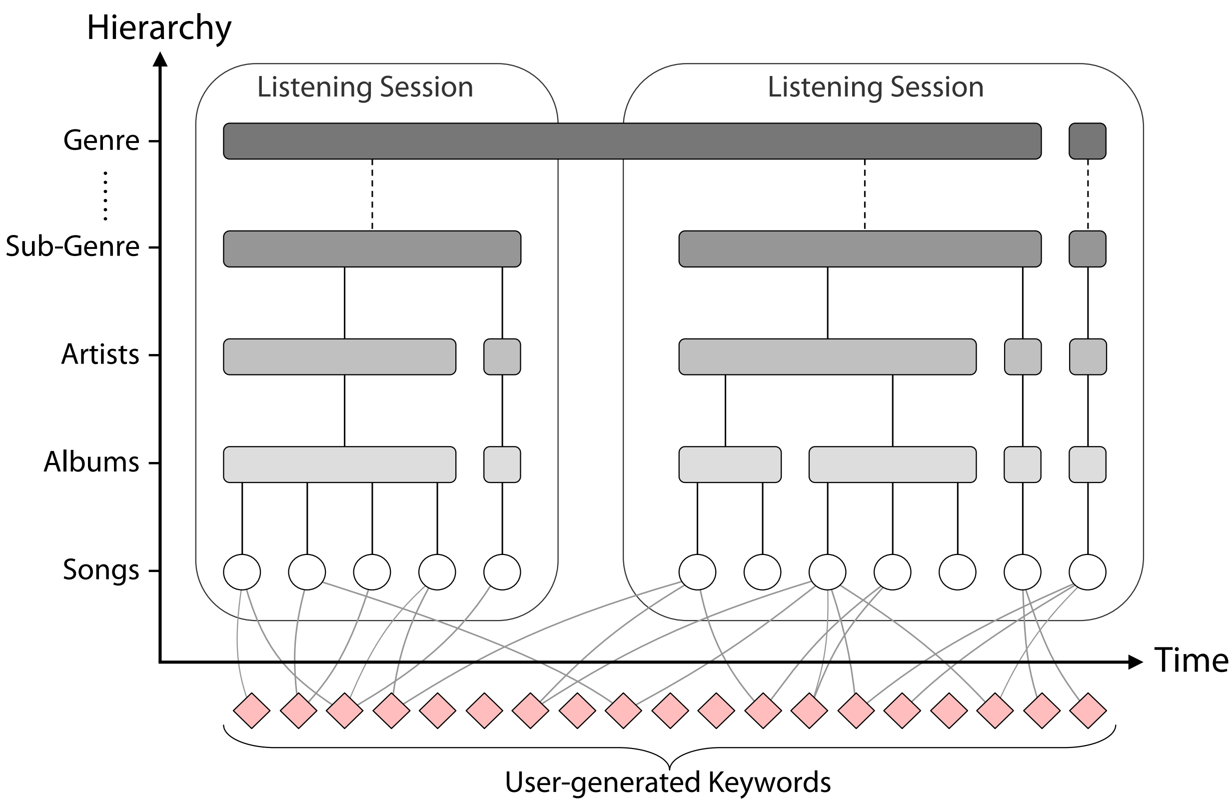

Figure 3.1 The space of (contemporary) music can be divided along the lines of single songs, albums, artists, and (sub-)genres. User-generated keywords or tags can provide additional information to the strict hierarchy (Source: [13]).

directed graph instead of such a tree would be problematic in the commonly used table-based media players (e.g., iTunes).

One way to overcome the restrictions of a hierarchy is to add user-generated keywords or tags. Last.fm lets listeners provide their own keywords for songs and also makes these tags available to the community. Using this wisdom of crowds [158], songs can be classified simply by taking the statistically most common keywords from the community. Regarding the contents of tags, Firan et al. identified 21,177 different keywords and found that: "... the top 100 most used tags showed that approximately 60% of the tags represent genre descriptions, while the rest of 40% is shared among tags describing personal impressions (e.g., 'seen live'), artists (e.g., 'female vocalists'), time period (e.g., '80s'), country of provenience, soundtrack, tempo, or instruments" [49]. The second category of personal impressions can also give insight into some of the lacking information from local and personal context. In a follow-up analysis, Bischoff et al. showed [18] that personal last.fm tags obeyed to the classification for user-generated keywords provided by Golder & Huberman [55]. This classification also contains categories for *self reference* ("seen live") and *usage context* ("workout", "study")[18] and can thus amend the available listening information.

User-generated keywords can, however, be no replacement for automatic collection of

contextual information: People might not be motivated to manually annotate all of their songs, leading to a semi-tagged collection that is too inaccurate to be useful. User-generated tags also have the usual problems of self-documentation and biases. Finally, due to their being costly to produce, tags are static and cannot reflect the ever-changing local context.
Last.fm not only provides tracks listening histories but also works as a social network. Profile owners can friend others and also style their profiles to reflect their personalities. This also means that last.fm collects and presents demographic information (gender, age, country of residence) for each of their registered listeners. Non-structured information can also be added in the 'About Me' textbox. All of this, however, is optional and prone to erroneous or fake entries.

3.3.3 Sources of noise and gaps

Even though last.fm's approach avoids the problems of manual data collection, there are still problems with the data that cannot be avoided. A (music) lifelogger's intention is to create a representation of the real-world events that is as accurate as possible. If the data is inaccurate there are either artificial memories that the actual owner of the lifelog never experienced or factual experiences have not been logged, which makes it harder to retrieve relevant information from such periods. A lifelog with too many blank spots can become useless and lead its owner to just abandon the undertaking. Also, when aggregating lifelog data missing information can skew the statistical or visualization results and render the overall data useless (e.g., when using a lifelog for dieting purposes and getting wrong calorie readings).
Such problems in lifelogs can be separated into two categories, namely *noise* (i.e., songs that have not been heard by the listener but appear in the data) and *gaps* (i.e., songs missing in the history) (cf., [13]).
Noise in an automatic listening history can come from several sources: "another person is using the same computer for listening to music and forgets to stop the scrobbling process or the user leaves the computer while the music continues playing"[13]. Also, switching off speakers or headphones can cause unhearable but scrobbling music playing. Noise in the data can also mean entries that are only in some regard erroneous: Wrong metadata and missing audio fingerprinting (e.g., in third-party player software scrobbling via the API) can cause the creation of listening entries where a wrong song or artist is shown. Software that caches listening history entries (e.g., when no network access is available) has to rely on a device's clock setting for the time stamp, which can of course be off. Entries that are directly transmitted carry correct time stamps as they are directly applied by last.fm's servers.
A very interesting source of noise in the lifelog is the listener her- or himself. As described above (see section 2.1.3), last.fm's public profiles also serve as badges of musical identity. In order to listen to embarrassing music or songs that do not fit into one's self image, the listener might decide to apply strategies such as diluting or

boosting [153] to cover up unwanted entries. Such an intentional skewing might even happen with private profiles, when their owners apply a self-directed form of identity management.
A similar sort of behavior can be caused by the listener's feeling of being observed. Last.fm's public profiles show every musical activity in real-time and depending on the listener's intention for self-presentation this can lead to her or him acting in a different way than in an unobserved state. The "public listening" introduces the other-directed identity construction use of music even to unaccompanied listening sessions. It can lead to last.fm profile owners listening to more or less music or choosing more deliberately than just by pressing shuffle, even when it is not clear if anybody accesses the profile. While this is not noise in the data processing sense, it still changes a person's listening behavior simply through the process of lifelogging (similar to the observer's paradox in physics).

A second category of errors in the data are *gaps*, where entries that should be in the log are missing. Reasons can be that "users may have a non-supported device for listening to music while on the go, they possibly use legacy media such as records or CDs, or visit a concert, a party or a club where music is played"[13]. Also, the aforementioned requirement for playing 30 seconds of a song before logging it can lead to blank spots. One major source of error are also problems with technology: Software bugs or missing network access can lead to failures in transmitting listening instances to last.fm. Again, listeners might also cause these problems themselves, for example through carelessness (forgetting to install or launch the Audioscrobbler or configuring a software player to scrobble). It is also possible to delete single entries or even a complete listening history in last.fm's web interface, which might be triggered intentionally (e.g., for a fresh start) or erroneously. Another reason might be the wish to be able to listen to music without being watched and switching the scrobbling actively off. Besides embarrassment this can also stem from paranoia, for example, when listening to pirated music (rumors that last.fm was giving listening data about not-yet released albums that had been leaked on the internet to record labels caused on uproar and corresponding cautiousness [27]).
A last set of gaps can come from last.fm themselves, when data is lost due to server malfunctions, hacker attacks or missing network access. Also, last.fm might be shut down one day due to economic considerations taking the data with it. While these reasons are imaginable they fortunately have not happened so far.
Various sources can cause noise and gaps in real-world listening histories and their owners have to work themselves to keep their online music representation as accurate as possible. Despite these problems, however, millions of listening histories exist on the service and one can assume them to be fairly accurate. These histories provide a fascinating look into the musical lives of many people.

[27] M. Arrington: Deny This, Last.fm, http://techcrunch.com/2009/05/22/deny-this-lastfm

3.4 A large-scale analysis of music listening

This availability of large numbers of listening histories also allows for more extensive, but at the same time accurate studies on listening behavior. Examples for this are Zheleva et al.'s study about modelling genre preference (based on 14 weeks of data for roughly 2,000 listeners)[178], Park et al.'s deriving of temporal patterns (with data from 70,000 people spanning almost one year)[127], and Herrera et al.'s similar approach with a smaller data set (992 people and four years of activity)[68].

We[28] were also interested in the expressiveness of listening histories (cf. [11]). As mentioned above, the main downside of real-world listening histories is their lack of contextual data, so we wanted to find out how much information can be drawn from last.fm's minimal listening histories when extended with other available music metadata. Research on musicology and human-music-interaction (see section 2) contained a large number of features that could be derived. One interesting feature is for example if listeners are interested in albums, which can be found using metadata and examining the order of listening instances. Our main goal of the study was therefore finding the most salient features of last.fm histories to gauge their expressiveness. The results of this study enable researchers and practitioners to estimate what features of a listening history are important to know and how they influence behavioral aspects: Would demographic factors determine the tendency towards album or playlist listening? Having this information allows improving recommendation or other customization activities using quickly collectable data. It also gives insight into the general behavior of last.fm listeners.

Information on the statistics of last.fm listening histories is scarce. While it is clear when the service had been founded, how old an average listening history is or how many songs it contains was less obvious. Therefore, our second goal became learning about the statistics behind the listening histories by taking a random sample and extrapolating towards the general listener base.

3.4.1 Methodology

We decided to obtain a representative data set from last.fm by using their API. The last.fm API allows accessing almost all parts of the service programmatically while also being generous regarding the terms of this access ("You will not make more than 5 requests per originating IP address per second"[29] is the most serious restriction). The API also provides access to demographic profile information (if available) and the complete listening histories.

Of course, this access is only available for public profiles: The owner of a profile is able to restrict access to her or his listening history. Yet, other profile information such as the

[28] This study was part of Jennifer Büttgen's Diplom thesis [26].

[29] http://www.last.fm/api/tos

profile image, user name, etc. is always visible. We made sure not to collect the real names of the profile owners to be able to perform our analysis anonymously.

Data set

For each listening history owner we collected the following overarching information:

- *username:* The profile owner's last.fm username.
- *demographic information:* The owner's country of residence, age, and gender (based on their profile's self-report).
- *subscriber:* Is the owner a last.fm subscriber and paying a monthly fee to support the service or not.
- *playcount:* The total number of tracks in the owner's listening history.
- *playlists:* The number of playlists the owner has created.
- *registry date:* When was the last.fm account created.
- *friends:* The number of last.fm friends the owner has.
- *shouts:* The number of shouts (i.e., publicly visible messages on the owner's profile page - similar to wall posts in Facebook) on the owner's profile.
- *private tracks:* Is access to the owner's listening history private or public.

In addition, we also extracted their listening histories (if they were accessible). Each entry of a listening history contained the following information:

- *music metadata:* Artist name, title name, album name, position on the album, duration of the track
- *streamable:* Is the track available for streaming on last.fm's webradio or directly on their website
- *toptag:* What is the most popular user-generated keyword describing this track

The first problem was finding a suitably random set of users. As last.fm does not provide user lists and manually guessing valid profile names would have been to costly, we decided to take another route: Last.fm's underlying recommendation is based on collaborative filtering that also allows calculating the similarity between two listening profiles. Similar profiles (so-called "neighbors") are made available through the API. We therefore decided to take a seed profile (one of the authors'), access its neighbor list, continue with the least similar neighbor (to provide a suitable diversity), and repeat

this process as long as desired. Informal tests showed that using this approach a small number of steps was enough to reach profiles with completely different tastes and demographic backgrounds.
A second problem was the API's terms of service's restriction to five requests per second. This meant that very large profiles (some examples had more than 370,000 tracks in their listening history [26]) would take a long time to download. We had, however, no data on the average size of a last.fm profile, which meant that defining a suitable cut-off point would be difficult. Therefore, we initially ran a script that traversed profiles using the above approach and only counted the total number of tracks in their histories. After 1,000 profiles we had an average profile length of 23,599.73 songs and decided to stop downloading profile data after 25,000 tracks [ibid]. For each such trimmed profile, we added a corresponding flag in the database.

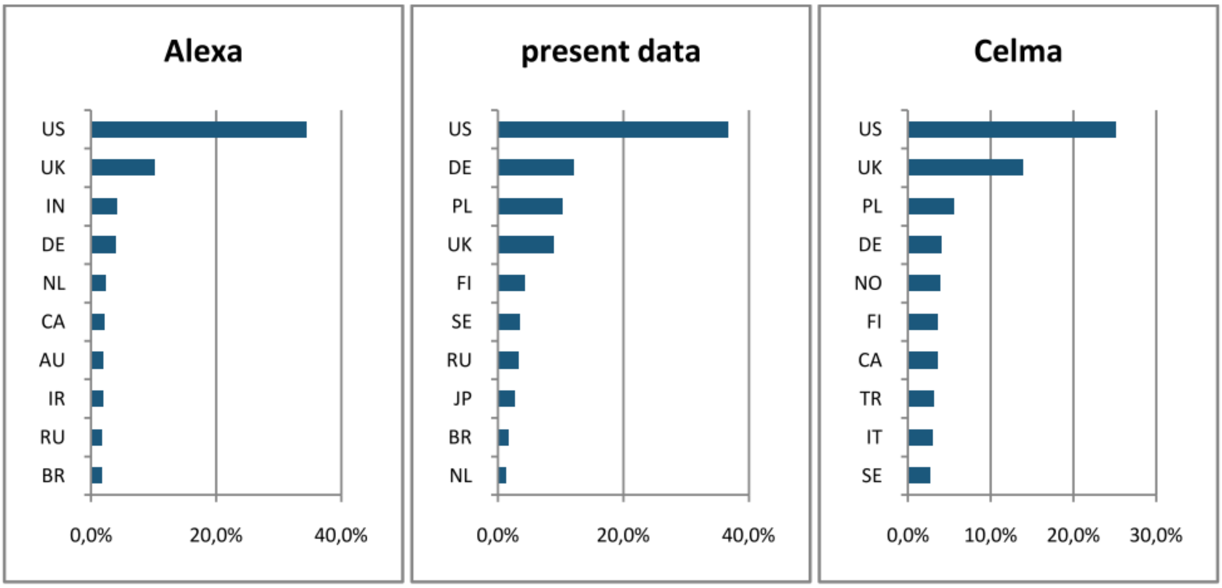

Figure 3.2 Country distribution of listeners in the Alexa, Celma and our data sets (Source: [26]).

Using this approach we collected 5,000 listening histories with over one million unique songs by 130,000 artists [ibid].
In order to make sure that the resulting sample was representative enough (and not biased by our selection mechanism), we compared the demographics of our set to two other sources about last.fm. One was by Alexa[30], a company that provides statistical information about website access. The other was a data set of 1,000 last.fm listening histories by Óscar Celma (see here[31] and chapter 3 of [30]). We extracted demographic information from our set and compared it with the information provided by these two

[30] http://www.alexa.com

[31] http://www.dtic.upf.edu/ ocelma/MusicRecommendationDataset/lastfm-1K.html

other sources.
A first point of comparison was the country distribution (see figure 3.2), which was similar across all three sets (USA is leading and Germany and the UK among the top four). Also, Celma and our set show a more similar distribution which presumably stems from the nature of the data sets: Both our's and Celma's are listening profiles while Alexa only tracks website access to last.fm.
Regarding age, Alexa only provides a comparison to the "general internet population", but acknowledges for last.fm a bias towards teenagers and people in their twenties and an underrepresentation of people beyond 35 years. Celma's data set had explicit ages and his profile owners' were on average 25.37 years old (SD: 8.315). Our distribution is corresponding, with an average age of 24.25 years (SD: 11.573). The gender distribution was less clear but still consistent: Alexa holds last.fm in line with the rest of the internet, i.e. slightly skewed towards male, and Celma's profiles were also predominantly owned by men (male: 50.6%, female: 38.51%, missing: 10.89%). Our data set showed the same tendency, but suffered from a more extreme distribution due to 83.78% of the profiles missing this information: The rest contained 11.36% male and 4.86% female profiles.
The overall similarity of these three sources led us to conclude that the data set was representative for last.fm despite our uncommon approach of selection.

Variables

Having downloaded and filtered a suitable number of listening histories, we started identifying suitable variables based on our available data and the literature. Variables could be either directly read from the data (e.g., number of friends) or calculated from the listening history (e.g., percentage of songs in the first half of the year).

Demographics: A first group of variables contained whether demographic information was available (*CountryAvailable*, *AgeAvailable*, *GenderAvailable*) and what other information had been collected from the general profile pages (*Subscriber*, *Playcount*, *Playlists*, *Friends*, *Shouts*). Also, *Active* held the number of seconds since the profile was created.

Temporal dynamics: We used a second set of variables to describe the changes over time. We planned to perform a principal component analysis (see below), so all variables had to be metric or binary. To keep the number of variables as low as possible, we opted for aggregating variables for single hours, months, etc. to overarching seasons and halves of the year. The resulting variables are *Balance* (fluctuation of plays throughout the year) and *yWinter* (percentage of plays in first half of the year). Another set of variables describes artists or genres that are only heard during one half of the year, one season, or day- or nighttime:
GenreUniques{Spr, Sum, Aut, Win, Halfyear1, Halfyear2, Day, Night} and
ArtistUniques{Spr, Sum, Aut, Win, Halfyear1, Halfyear2, Day, Night}.

General listening behavior: A last group of variables describes the overall listening behavior: *TracksDay* (Average number of songs per day played), *SongLg* (Average song length), *SessionLg* (Average session length in songs). The distinctiveness of a history is determined by dividing the number of unique songs, albums, or artists by the play-count: *Distinct{Sng, Alb, Art}*. The next group of variables describes listening sequences: if albums or an artist's albums are heard in the right track- or chronological order (*InOrderAlb, InOrderArt*) and completely (*ComplAlb*), and if songs are skipped (*Skips*) or listened to in shuffle (*ShuffleAlb, ShuffleArt*). Interesting are also direct repetitions of songs, albums, and artists (*Perma{Sng, Alb, Art}*), sessions that start with the same song (*Introsong*), or the percentage of songs that are only listened to once (*Onetimer*).
One problem we faced was that last.fm's webradio also scrobbles, so our analysis might end up with a description of their underlying algorithm instead of a person's listening behavior. Therefore, each sessions where ten songs in a row were streamable was flagged with *Online*. Finally, we also had a variable describing the discovery of an artist through a newer album of theirs (*DiscoverArt*) and the changes in genre and artist taste from one year to the next (*LoyalGenre, LoyalArtist*).

In order to make the results of the analysis as representative as possible we had to make sure that the used listening profiles represented a stable and reliable version of a person's taste in music. Taking profiles that had only been created recently and whose owners were still in an experimentation phase regarding the service could have skewed the results too much.
Also, we had several variables that were geared towards longer stretches of time (all variables concerning seasons or halves of the year) and the two *Loyalty* variables that explicitly compared data from one year to data from the next. We therefore filtered our initial 5,000 histories and only kept histories that had been active for at least two years which left us with a set of 310.

Analysis

After deciding for 48 variables, we had to analyze their actual impact on the profiles. Our selection of variables had been taken from the literature and personal experience, but it was not clear how important or representative they would be for a set of different listeners. Given this flood of variables, we wanted to learn what the most distinct descriptors of a person's last.fm profile were and which of the variables were most important for distinguishing between these histories. This way, simple analyses of short listening sequences or, in the best case, just having some demographical data can provide a wealth of information about the person's listening preferences.
A common approach for reducing high-dimensional data sets to its most distinctive features are dimensionality reduction techniques. We opted for a Principal Components Analysis to make sure that the number of variables was reduced which would not be guaranteed using other related techniques (Independent Component Analysis possibly even increases the number of dimensions, while Exploratory Factor Analysis only identifies common variance). Also, as this was the first analysis looking at a broader picture

of music listening in automatic histories (compared to only temporal dynamics [127, 68] or genre taste [178]), PCA was chosen to produce clear results and form the basis for other, similar studies.

3.4.2 Results

Our analysis produced results in two areas: General information about last.fm profiles, and the impact of various factors (encoded in our variables) on the distinctiveness of these profiles.

Statistics of last.fm profiles

The first set of results were general statistics, so we used all 5,000 profiles for this analysis. We were first interested in the willingness to share information, i.e., privacy concerns and the willingness to use last.fm as a representation for one's taste in music. Interestingly, regarding demographic information (age, gender, country), people appear to think on an 'all-or-nothing' basis: 82.74% did not provide any information, 1.93% and 4.24% provided one or two pieces, and 11.09% all three. Maybe the inhibition to fill in the corresponding text fields lowers once somebody has started doing it (similarly to the skipping of songs).
People were less reluctant to share the listening histories themselves: 99.08% of all histories are publicly visible, which might stem from the general anonymity on the platform. Financial support for last.fm through a subscription is provided by 2.40% of all profile owners.
Last.fm also works as a social network for music: Each profile owners has on average 5.45 friends (SD: 127.075) and 6.73 shouts (SD: 51.88) left on her or his wall. Yet, the rather high standard deviations show two different approaches to using last.fm: either for listening to music or as a social network.

Regarding the overall listening behavior we were also able to extract interesting information from the data set: If last.fm profile owners were generally stable and similar in their listening behavior, there should be a correlation between the total size of their listening histories and their registration dates (i.e., the amount of time they had to fill their profiles). The results, however, show that there is no correlation between these factors (see figure 3.3) and the play count is much more dependent on the involvement with music than the time on the service: The mean listening history had a size of 4260 tracks with a standard deviation of 11822.197. Analyzing the overall time of day when music was consumed showed an expected distribution (see figure 3.4): Similar to North et al.'s study [121], we found peaks in the evening hours and a general decline throughout the night, reaching its lowest point at 6AM in the morning. This analysis contained a sample of 100,000 songs from listeners who had provided country information, and is surprisingly stable worldwide: Simi-

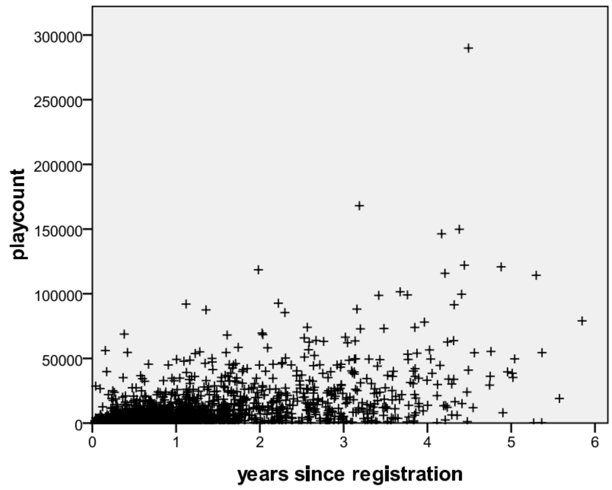

Figure 3.3 The size of a listening history compared to the years of activity on last.fm shows no correlation (Source: [26]).

lar analyses for various countries in the data set showed identical distributions (cf. [26]).

Most salient factors in last.fm profiles

A second analysis concerned the principal components in the 310 suitable last.fm profiles and thus the importance of single variables for distinguishing listening behavior.
We first tested if the variables in the data were uncorrelated using Bartlett's test of sphericity which was supported by the results (χ^2= 13425.536, df = 1128, p < .001). In order to make sure that the correlation matrix could be used for the factor analysis, we also performed Kaiser's Measure of psychometric Sampling Adequacy, which resulted in a value of 0.825 (meaning suitable for the analysis). The lowest communality of a variable in our sample and the only one below 0.5 was *SongLg* with 0.476 which we found still acceptable.
After these preliminary tests to make sure that our sample is suitable for a factor analysis we determined the number of components using Kaiser's criterion and Cattell's scree test. After comparing preliminary results with four (Cattell) and 13 (Kaiser) components, we decided for Kaiser's 13 components, which contained 74.278% of the total variance. To extract the 13 components, we rotated the component matrix using the VARIMAX algorithm (the resulting matrix can be found in the appendix C). The resulting 13 factors (see table 3.1) describe interesting connections between the variables.
One major influence on the resulting histories are temporal cycles: Factors 1, 3 and 7, that together explain 37.8% of variance, include all variables for unique genres and artists for the four seasons, the halves of the years and day and night. This clearly shows the impact of seasons on music listening and how people's affectedness to it can be used to distinguish their listening behaviors.

Nr.	Name	Variables	Var.	Description
1	Distinct Music 2nd half-year and at night	*ArtistUniques {Sum, Aut, Halfyear2, Day, Night}, GenreUniques {Sum, Aut, Halfyear2, Night}, DistinctArt, Onetimer*	**27.3%**	Unique genres and artists listened to in summer and autumn and during nighttime. It also contains the distinctiveness of artists and the tendency to hear songs only once.
2	Unique versus repeated music	*Perma {Sng, Alb, Art}, Distinct {Sng, Alb}, ComplAlb, Balance*	**8.3%**	Tendency to repeat versus to listen only once, also to complete albums. Finally, the steadiness of listening throughout the year.
3	Distinct Music 1st half-year	*ArtistUniques {Spr, Halfyear1}, Genre-Uniques {Spr, Halfyear1}, yWinter*	**7.2%**	Unique genres and artists listened to in winter and spring. Also, the tendency to listen to more music during the first half of the year.
4	Complete profile	*CountryAvailable, GenderAvailable, AgeAvailable*	**4.9%**	Willingness to fill out demographic information in the profile.
5	Listening activity	*Playcount, TracksDay, SessionLg, ShuffleArt, Introsong*	**4.2%**	Listening intensity and duration, shuffle listening to artists, the tendency to have intro songs.
6	Stability	*LoyalGenre, LoyalArtist*	**3.7%**	The stability of musical taste from year to year.
7	Winter listening	*ArtistUniquesWin, GenreUniques {Win, Day}*	**3.3%**	Tendency to listen to special genres or artists in winter.
8	Album listening	*ShuffleAlb, DiscoverArt*	**2.9%**	If albums are shuffled and artists discovered through them.
9	Webradio	*Subscriber, Online*	**2.7%**	Subscriber status and tendency for webradio listening.
10	Activeness	*Active, SongLg*	**2.6%**	Age of profile and average song length.
11	Sequentiality	*Playlists, InOrderAlb*	**2.5%**	Tendency to listen to albums in order or create playlists.
12	Social Networking	*Friends, Shouts*	**2.3%**	Use of last.fm as a social network.
13	Through Artist Skipping	*Skips, InOrderArt*	**2.1%**	Skips through albums and listening to artist albums in order.

Table 3.1 13 principal components describing last.fm profiles (Source: [26]).

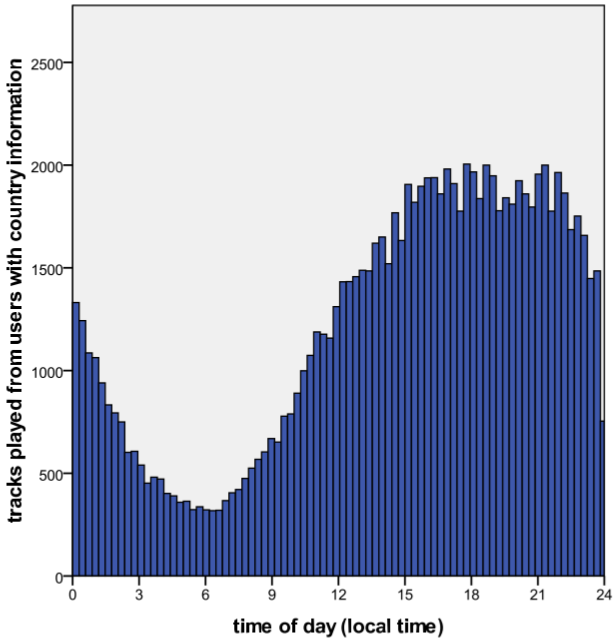

Figure 3.4 The distribution of music throughout the day shows peaks in the evening (Source: [26]).

A second set of factors derive from a person's song-to-song listening behavior: Especially factor 2 (8.3%) is interesting, as it includes somebody's tendency to repeat songs, albums or artists and also the opposite factors of having many distinct songs and albums in one's collection. There exists a clear dichotomy between listeners who enjoy an album repeated times in a row (the variable *ComplAlb* also has a positive factor loading for this component) and those who only listen to songs once and this difference is also visible in the variables. Similarly interesting is factor 5 (4.2%) that contains the intensity of listening (number of songs per day, length of listening sessions, etc.).
Stability (factor 6, 3.7%) contains the appearance of both genres and artists in subsequent years and also had an influence on the rest of the history. In contrast and as visible from the general statistical analysis, the age of the profile had only minor influences on the rest of the listening behavior (factor 10, 2.6%).
Finally, there are also factors regarding the use of last.fm as an online platform, where factor 4 is surprisingly expressive (4.9%). Knowing that the owner completed relevant demographic information can therefore also be a distinguishing factor for the rest of the profile (possibly due to the identity construction function). Factor 12 (using last.fm as a social network), however, is only of minor (2.3%) importance.

Correlations between age, gender, and listening

One goal we had with this study was to provide a simple way to predict listening preferences through minimal questioning of the listener. Demographic data would of course

be perfect for this, as minimal input could already provide large benefits. The PCA data set only contained information if demographic information was available but not which, so we performed another analysis on another subset. This subset had 546 profiles that contained age and gender, but were not necessarily two years old, so we had to drop all variables that required this (all *ArtistUniques*, all *GenreUniques*, all *Loyal*, *Balance*, and *yWinter*). We had an overlap of 196 profiles between the two subsets. We similarly extracted correlations from the data (cf. [26]), and looked at the influence of *Age* and *Gender* on the results:
Correlations we found between *Age* and other variables included a positive one with *Distinct {Sng, Alb, Art}* (r(544) = +.314/+.341/+.278, $p \leq 0.01$) and *Onetimers* (r(544) = +.318, $p \leq 0.01$), which means that older listeners have a more diverse taste in music. They also enjoy more *Webradio* (r(544) = +.259, $p \leq 0.01$) and listen to more music in the morning (r(544) = +.301, $p \leq 0.01$), probably before going to work (there is also a slight negative correlation with listening on the weekend (r(544) = -.191, $p \leq 0.01$)).
The differences between the two *Genders* were less expressive: Male listener's profiles were slightly older (r(544) = +.126, $p \leq 0.01$), and had a higher tendency towards using the same *Introsongs* (r(544) = +.166, $p \leq 0.01$). Our conclusion is that in terms of listening behavior, both men and women act similarly and more or less involved listeners exist on both sides. Looking at the genres and artists in detail might have unearthed more striking differences but that was beyond the scope of this study.

Discussion of the results

This study is the first of its kind that includes all available factors of automatic listening histories and subsequent analyses to arrive at a description of which variables are important for distinguishing between listeners, independent from specific genres. The results describe first of all the relationships between the different variables and which ones can be used as stand-ins for the others (all members of one component). Second, they also provide an order of importance for each of these factors, showing the impact of seasonal and cyclical behavior and repetitions and single plays in collections. The factors allow separating the community into groups and stereotypes which can be used for recommendation and also for automatically adapting interfaces: Depending on the current season, different songs might be recommended. Listeners who have a tendency to repeat songs might get suitable automatic playlists and will be presented with a prominent option to repeat the current song. Distinct listeners on the other hand can receive more unusual playlists that are more suitable for discovery. The interesting correlations between demographics and different variables can also be used to categorize listeners simply by asking them for their age and then cater to their respective needs.
One caveat of the results is that last.fm profiles are very specific representations of listening behavior with their own little quirks. Noise & gaps in the histories (see section 3.3.3) are always prone to skew the results but should disappear due to the size of the sample. Also, detecting skips that happened in the first 30 seconds of a song was not possible, as the Audioscrobbler simply ignores them. The roughly one percent of private

profiles might have hidden an interesting group of listeners. And finally, scrobbles from the webradio (in the *Online* variable) are also problematic but had not a large influence on the results as the low impact on the factor analysis (2.7%) showed.
Beyond its immediate results this study is also relevant on a meta-level. As described in section 2.3.2, musicologists and sociologists have sophisticated approaches for collecting real-world music listening data, but some of the burdens might be taken from them and their participants alike simply by relying on automatic data collection. Instead of letting participants note what song they are currently hearing at several times during the day, researchers could simply give them a smartphone that collects audio information for identifying the songs afterwards. This in turn would also increase the accuracy from one or several samples a day to complete song-to-song listening histories. For contextual information researchers can still rely on diary entries or questionnaires.
Also, if context is not central, researchers can rely on last.fm to provide a listening data set containing histories for millions of people. For some studies it might no longer be necessary to collect data, but to simply form a hypothesis, take a suitable sample from last.fm, apply some statistics and immediately confirm or reject it. Lifelogging data has the potential to revolutionize research in the humanities.

3.5 What's in a listening history?

Listening histories are an interesting type of lifelogging data, in that they are relatively easy to collect and there already exist large data sets. Compared to other lifelog movements that are (currently) more niche, they show the feasibility and benefits of lifelogging even for regular music fans. They are, however, certainly different from their ideal versions. In this section, I will describe what real-world histories can and cannot tell and how to make up for their shortcomings.

3.5.1 Mapping real-world data to the ideal

As we've seen, there are differences between ideal listening histories and their real-world counterparts especially regarding the inclusion of contextual information. Last.fm's representative approach to collecting listening histories comes down to the basics of linking a listener to a song and a time stamp. Compared to our hypothetical ideal listening history that was based on the various listening factors, they fall short in several regards:
Regarding the *Personal Context* (see section 3.2.1), last.fm provides almost nothing beyond very basic demographic information such as age, gender and country (if at all). The informal "About me" part of the profile might contain information about personality, intellectual capabilities, musical training, or access to streaming services and mobile players but this is rather improbable. Other aspects such as musical preferences or the

size of the personal collection, however, can be calculated just by looking at the collected history.
The situation is better for the *Local Context* (see section 3.2.2), as both date and time and the musical context are either directly collected or can be accessed based on the unique identifier of the song. Our own study (see above) has shown what amount of variables can be produced by identifying behaviors (e.g., listening to the same song repeatedly) and processing the data accordingly. Other local aspects, however, are completely unavailable: The emotional state of the listener just like her or his current task and social surroundings are valuable information for explaining listening decisions and unfortunately lacking. Collecting them, however, would require a far greater effort by last.fm and probably also by the owner of the lifelog to be feasible. Some of this context might be derived from user-generated keywords, but they, as mentioned above, require involvement and are static.
The *Global Context* (see section 3.2.3), finally, might be relatively stable throughout the listener population but lies (mostly) beyond last.fm's grasp. The impact of record labels on music and relevant technological developments are not available from the Audioscrobbler, just as the role of the social acceptance of music. Yet, last.fm can make statements about current musical trends based on their massive data base of listening histories. They also provide this information via their API and create charts based on it, similar to the existing sales charts (e.g., Billboard). The accumulated listening histories can therefore provide some global context.
Despite their shortcomings, real-world listening histories therefore represent a valuable type of lifelogging data: They not only stand as a proof-of-concept and teach non-lifelog-enthusiasts about possible benefits, their sheer number also provides a backdrop for individual ones (compared to the "silos" of other proprietary lifelogging systems). Also, the lack of data can be compensated by relying on other available data sources for context.

3.5.2 Additional data sources

In order to create a perfect lifelog, a hypothetical solution might simply take a mobile device and add all necessary (even more hypothetical) sensors for a perfect, around-the-clock tracking. However, most of lacking aspects of ideal listening histories are far too complex to be captured with a sensor (e.g., the social surroundings and their current implications), or not available for sensing (e.g., the listener's emotional situation). Yet, even without such a perfect, unattainable sensing device, at least some of the missing context can be restored through other sources: One common dimension of all lifelogging data is time and this attribute also allows merging data from various sources. This section contains possible data sources for enriching listening histories from three areas: *Episodic memory*, *Local data*, and *Online services*.

Episodic memory

One powerful data source that sometimes allows for explaining otherwise cryptic listening decisions is the listener's memory. Lifelogs can only serve as extensions to biological memory as they lack data and contain mostly low-level information (e.g., blurry photos, non-descript audio recordings, single songs). In order to make sense of this data and tell the stories that happened at that time, the owner of the lifelog is needed (for this section cf. [13]).

It is interesting to compare the attributes of lifelogs and biological memory: While the former contains perfect representations of data with little significance, the latter stores vague renditions of events condensed to their main aspects (if needed, details can always be re-imagined on-the-fly as so-called *false memories*[141]). The former also allows direct access to any data based on its time stamp, while the latter struggles with exact dates ("What did you have for dinner last Tuesday?"[13]) but provides rich, content-based associations between items.

One main attribute of human memory is that its store of autobiographical events is organized episodically [162]. Such episodes can be accessed using memory triggers that can be of very different nature and based on all sensory impressions: from visual cues, via melodies, voices, to smells or tastes (Proust's narrator in *Á la recherche du temps perdu* starts his seven-volume narration based on the memories triggered by the taste of a madeleine). Human memory, however, is also highly efficient storage-wise and only keeps information that it deems relevant at the time. Therefore, even providing "right" memory triggers might not cause a recall, simply because the person's mind at the time decided to ignore this fact. A way to solve that is to store as much information as possible electronically to be able to provide a wealth of possible triggers.

A second aspect that is important regarding episodic memory is the so-called *generation effect*[82]: People are better at remembering items created by themselves (written text, photos), than ones they only consumed (read text, seen photos) and this even works (partially) with automatically captured ones (Sellen et al. showed this in a study about this aspect with Sensecam pictures [148]). Therefore, episodic memory triggers that allow understanding a lifelog should be as varied as possible and created by the owner of the lifelog, if possible.

Local data

Various types of data are created regularly by people using computers (electronic communication, written pieces) or stored on them (photos, images) and can be made available for enriching lifelogs. And there is also data that is actually used for organizational purposes but also contains context, most prominently calendars.

Lifelog applications can simply rely on this data as memory triggers. *Photos* are one of the most powerful sources for not only associations but also context. Oftentimes, people only have to see a glimpse of a photo to remember the surrounding event and their feelings at the time. Also, the function to take pictures is by now included in almost any portable device from cell phones to media players, so that the problem of creating

photos has long been displaced by the problem of finding and organizing them.
Another important source are *calendar entries*. While they are mainly used as prospective devices for remembering and planning, in hindsight they also become parts of and enrichments for lifelogs. Depending on the details in an event description they can contain information about current tasks, problems and also social surroundings.
Similarly, (locally stored) *electronic communication* such as emails or instant messaging can also work as a backdrop for episodic memory. Highlighting relevant topics of a certain time [166] or the most important contacts can make other decisions easier to understand.
In the context of listening histories, actual representations of the *music* heard are valuable to interpret the logs. Even complex textual descriptions with extensive meta-data cannot capture the richness of music, so simply playing back a certain piece from a history is often the easiest and most expressive way to help the owner remember.

Online services

Not every aspect of one's life is kept on the personal harddrive and more and more tasks and tools are transferred to online counterparts. Photos are moved to flickr and Facebook, and calendar entries and electronic communication to Google Calendar or Mail, or Facebook. The trend towards cloud-based computing promises to rather sooner than later move every type of personal data to online databases coupled with easy synchronization and access. The downsides regarding network failures and privacy only show little effect on the growth rates of online services.
In addition to transferring existing data to the cloud, online services also capture other, formerly non-existent data: *Status updates* in social networks such as Facebook, Twitter, or Google+ make the idea of a static "About me" representation dynamic. Facebook's ("What's on your mind?") and Twitter's ("What's happening?") titles for their status boxes both aim at learning about the current context in an informal way. In connection with a lifelog owner's memories such status updates can be very helpful in understanding.
Another automatically or semi-automatically collected source of expressive data are *location histories*. Services for smartphones such as foursquare[32] or Gowalla[33] capture (manual) "check-ins" to different venues. Automatic versions such as Google's Latitude[34] take location snapshots in regular intervals. Location data is helpful in retrieving specific information ("Where have I gone for lunch yesterday") but also allows deriving context: Specific venues stand for specific activities, and such information can be as helpful as manually-created calendar events.
Finally, music listening histories can be enriched with all the metadata available online. The fact that songs in such a history are uniquely identifiable allows adding informa-

[32] http://foursquare.com

[33] http://gowalla.com

[34] http://www.google.com/latitude

tion about its content (using an analysis service such as The Echo Nest[35]), its current trendiness (provided by last.fm and also Echo Nest), or other information such as the biographical background of the artists (via Musicbrainz[36] or Wikipedia).

3.6 Summary

In this chapter, I discussed the concept of lifelogging and music listening histories as one concrete instance of this idea. Lifelogging projects are currently focussed on the problems of capturing and storing the data, while accessing it is not yet central. Afterwards, I presented the attributes of an ideal listening history and put it in contrast to its real-world implementations that suffer from lower complexity and noise and gaps. I then gave details on a study that we performed and where we analyzed 310 long-term listening histories in detail. The results showed that seasons and personal taste were very important for listening decisions and that a surprising level of descriptive detail can be reached by calculating aspects and enriching the data with additional meta-data. In the last section of the chapter, I discussed what real-world listening histories lack and how to make up for it using additional data sources.

[35] http://the.echonest.com

[36] http://musicbrainz.org

Chapter 4

Visualizing Listening Histories

This is one for the good days
And I have it all here
In red, blue, green
Red, blue, green.

– Radiohead - *Videotape* –

Music listening histories are a large and complex type of data. As we have seen in the previous chapter, an average history contains thousands of unique tracks and can span several years. Providing listeners with simple lists of songs and time stamps as last.fm does might not be enough to allow understanding underlying patterns.

To provide this understanding, having a *summarizing* approach to presentation (cf. last chapter 3.1.2) for the listening history can unearth relevant patterns that were not obvious before. Simply calculating charts or trends from the data shows the most favorite artists, genres and songs and allows matching the taste between friends (last.fm presents this comparison in taste as a simple, one-dimensional *musical compatibility bar*, but at least explains which artists overlap in both profiles). Yet, summarizing approaches always suffer from the inherent abstraction in their presentation. Digging deeper and trying to understand how a certain result came about (e.g., why a musical compatibility value is so low) is usually not possible - the source and explanation for all such summarizing is the black box of statistics.

Yet, there is no way around summarizing, as having direct access to all items equally works only up to a point. With thousands of songs and artists accessing each one manually is not feasible and does not contribute to seeing a bigger picture.

One methodology for solving this problem of understanding and exploring is interactive information visualization. Information visualization techniques rely on the massive

parallelism of the human optical system to allow for a more effective transmission of information. While text-based displays of information force attention to hop from item to item, visualization displays all relevant information at once and allows a self-chosen path of exploring the data. Combined with interaction capabilities to filter for relevant items or adjust the display based on some chosen parameter (e.g., sort by an attribute), visualizations are powerful tools for exploring and understanding.
With all the strengths of visualizations comes also a certain complexity. Being able to use a visualization tool in detail and knowing about all its aspects can require dedicated training. Additionally, depending on the data being visualized, understanding all traits of the data might even need a solid background in statistics. So, while visualization is a valuable tool for making data graspable, it can be overwhelming for non-visualization-experts. The imposing rows of buttons and interface widgets of standard tools like Tableau[1] can already signal to the uninitiated that this software was off limits, just as the resulting graphics and textual output. Yet, the owners and producers of lifelogs are usually also the most skilled experts on the data, so their analysis can produce the most insightful results.
A solution to this dilemma between complex software and complex data is casual information visualization [133]. The complexity of a visualization interface can be reduced by focusing on important or interesting questions instead of trying to support every query imaginable (and additionally allowing for discovery of patterns in the data that one was unaware of before). Pousman et al. explain the difference between traditional infovis and casual infovis: "In core infovis, a system should have a tool-like ability to do work to display data, uncover trends and outliers, and generate hypotheses. Casual Infovis systems are useful artifacts that are helpful for providing representations of data, but without a clear task focus" [ibid]. While Pousman et al. also included more abstract cases such as artistic visualizations (that often completely ignore the exploration and are very explicit in their statements) into their casual terminology, I argue that even casual visualizations can produce serious insight. An aesthetically pleasing appearance and a presentation that emphasizes interesting bits about the owner's data can entice even non-infovis-experts to interact with a visualization tool. Once this initial stumbling block is cleared, people are ready to dive deeper into the analysis and play with their data, given the tool remains friendly and understandable. Finally, relying on existing and known interface metaphors and keeping the interface itself as simple and non-threatening as possible can allow even non-experts to analyze their complex lifelog information.
In this chapter, I will first present the previous work on the visualization of listening histories. These fall into two categories: *summarizing* variants that present an abstracted overview and *single-purpose* versions that succinctly describe one interesting aspect of the data without allowing much exploration. The main part of the chapter is taken up with the discussion of the design space of visualizations for music listening histories and its main design dimensions and additional aspects that are useful to keep in mind.

[1] http://www.tableausoftware.com/

4.1 Related Work - Visualizations for Listening Histories

The convenient availability of last.fm listening histories through their API made their data a popular target for simple visualizations and statistical tools. Some of the profile pages are embroidered with various widgets and applications that stand as proof for the enthusiasm of their owners and their adorement of their musical idols. They are often used for putting an explicit emphasis on a certain fact about one's listening behavior (e.g., 'Top-10 listener of Justin Bieber').
Accordingly, most visualizations that are produced are single-purpose or an a very abstract level. They are good enough for a short aha moment, but are not used beyond that. Usually, these small visualizations do not allow exploration and are not even interactive. The approach is that not all tasks are supported in a single tool but spread across the landscape of available last.fm visualizations.
Almost all of the visualizations available for listening histories have been written by fans or companies. Only rarely is research involved with this type of data, even though visualization of lifelogging data can be expected to be relevant in the near future. Right now, all this data is collected without suitable ways for accessing it.
Visualizations for listening histories can be separated into summarizing and single-purpose approaches.

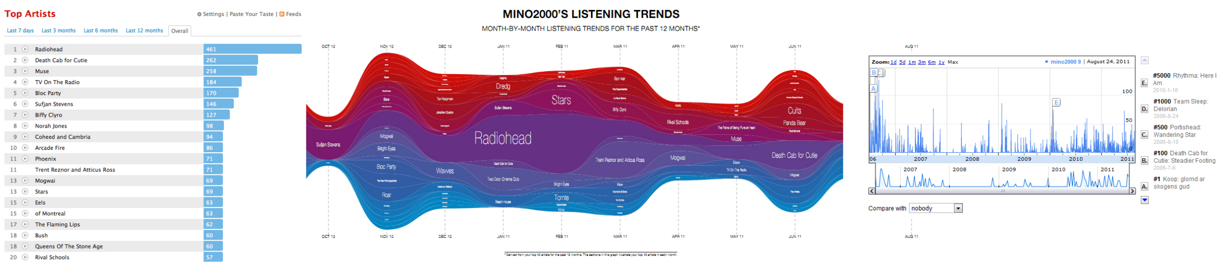

Figure 4.1 Three examples for summarizing listening history visualizations: (a) bar charts for top artists, (b) streamgraph, (c) histogram for scrobbles (Source: last.fm).

4.1.1 Summarizing visualizations

Summarizing or abstract visualization provide a view of a listening history on a very high level. They usually focus on longer periods of time and provide an overview, yet do not allow to explore the data in detail.
A simple example of such visualizations are the diagrams presented by last.fm themselves on the profile pages (see figure 4.1 (a)): A listener's favorite artists and tracks

are automatically summarized as bar charts. Normalisr[2] provides similar functionality, but relies on the actual song duration (instead of taking each song as having the same length) and optionally produces a nicely layed-out map of the top artists' photos. More charts are produced by My Music Habits[3] (maps the tendency towards discovery of a listener, i.e., the ratio between known and unknown artists in a history), Last.fm Tools[4] (tag clouds and statistics about neighbors and friends), or in rather abstract form in last.fm's Music Universe[5] (music taste as planets in various sizes and shades).
A large group of visualizations uses a time-line as basis for their presentation. A well-known example are Byron & Wattenberg's Stacked- or Streamgraphs [27] that give an abstract and aesthetically pleasing overview over a set of time-based data (see figure 4.1 (b)). The authors also applied the technique to a last.fm history on the level of artists. Streamgraphs are very effective in displaying time-bound data from up to several dozen categories, but quickly lose their appeal when too many categories shrink labels into invisibility. Therefore, some abstraction is always required and visualization on the level of songs is not sensible. Due to their appeal and effectiveness, however, streamgraphs are commonly used in visualization and last.fm alone has three tools for creating them (Listening Trends[6], lastgraph[7], Last.fm Explorer[8]). Other abstract time-centric visualizations are last.fm's Tube Tags[9] (prominent genres shown as a pseudo-London underground map) and Sha Hwang's Last.fm Spiral[10] (spiral visualization of top artists) that both highlight artist preference over longer timespans (weeks). A different way to look at artists' influence over time is the History Charts visualization[11] that draws histograms for each of the top last.fm artists of a listener over time. Finally, last.fm also provides a histogram for all scrobbling activity with the Scrobbling Timeline[12] (see figure 4.1 (c)) that invites day-to-day exploration of song entries, and Explorer[13] creates various bar-charts and histograms summarizing listening activity over time. A more recent example are the Last.fm Heatmap Calendars by Martin Dittus[14] that display scrobbling intensity on a daily basis. These various visualizations provide an overarching picture of one listener's activity, from preferred genres and artists to general scrobbling intensity.
Another set of visualizations relies on last.fm's global data collection to produce overar-

[2] http://www.normalisr.com

[3] http://www.mymusichabits.com

[4] http://lastfm.dontdrinkandroot.net

[5] http://playground.last.fm/demo/musicuniverse

[6] http://playground.last.fm/demo/listeningtrends

[7] http://lastgraph.aeracode.org

[8] http://alex.turnlav.net/last_fm_explorer/

[9] http://playground.last.fm/demo/tagstube

[10] http://www.diametunim.com/muse/

[11] http://playground.last.fm/demo/historychart

[12] http://playground.last.fm/demo/timeline

[13] http://twothreefall.co.uk/lastfmexplorer

[14] http://www.flickr.com/photos/dekstop/sets/72157627536754640/

ching impressions not for single listeners but larger groups of people or even countries. Adjei and Holland-Cunz show several approaches for making overarching patterns visible in their Diplom thesis [1]: comparing taste in genres of specific listeners, displaying the fluctuation of fans based on countries, uncovering the impact of an album after its release and mapping performances of artists from one genre to their locations. The resulting charts are beautiful and show trends beyond single listeners, yet they are not directed towards a single person (they allow no comparison to one's own situation) and are usually not interactive and explorable. Similarly, World Chart[15] displays the popularity of an artist world-wide based on all last.fm listening data.

Summarizing visualizations commonly allow integrating a second listening profile into the graph for comparison: Last.fm Tools derives statistics from neighbors, Listening Trends and Scrobbling Timeline show entries for selected friends on the same timeline, and Pretzlav's Last.fm Explorer integrates other profiles into the streamgraph. Some summarizing visualizations are also explicitly built for comparison. Chen et al.'s HisFlocks[33], for example, displays color-coded artist labels for two profiles over time and maps their spatial positions to overarching genres. By adjusting the visible time frame week-wise, the fluctuation and importance in artist taste can be shown. Hao Su's Last.fm Similarity[16] draws circular barcharts that highlight artist popularity for one or more listeners.

All of these summarizing visualizations have in common that they provide an overview over all or large parts of the available data. Looking at them provides a rough, statistical image of a person's taste in music or listening behavior. To learn more about the details or find an explanation for an uncommon result requires other tools.

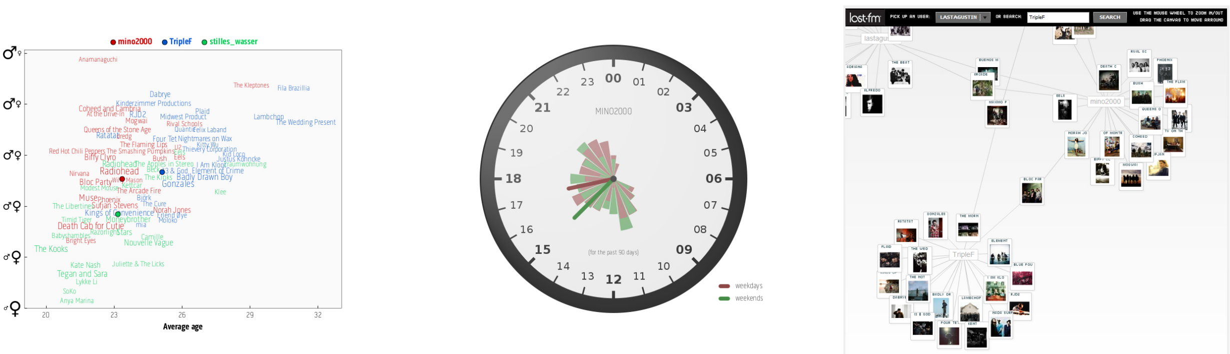

Figure 4.2 Three examples for single-purpose last.fm visualizations: (a) favorite artists mapped against average gender and listener age, (b) listening over 24 hours, (c) a listener's social network (Sources for (a) and (b): last.fm, for (c): Agustin Abreu).

[15] http://playground.last.fm/demo/worldchart

[16] http://www.haosu.org/2010/10/31/last-fm-similarity

4.1.2 Single-purpose visualizations

A different way to approach listening histories are single-purpose visualizations. These focus specifically on one of the details and on answering a single question concisely based on the listening history. These answers might still be vague regarding the underlying songs (e.g., by simply aggregating their numbers or working on an artist-level) but tell a different story than general taste or temporal aspects.
Various types of badges are available for expressing one's allegiance to artists (e.g., by listing logos of one's top-10 bands[17]), albums (e.g., with recently listened album covers[18]), or one's amount of scrobbled tracks (e.g., as an abstract "last.fm level"[19]). These badges are small pieces of HTML that can be added to the 'About Me' box of a last.fm profile and some such pages are cluttered with them.
More sophisticated examples for single-purpose visualizations use the last.fm community as backdrop for one's own taste in music. The Underground-O-Meter[20] separates profiles into "mainstream" and "underground" based on the number of fans of the top artists. Last.fm's own Gender Plot[21] (see figure 4.2 (a)) takes one or several listeners' top artists and plots them against the average age and gender for that artist. The result is a "last.fm age" and "gender" with explanations in the form of the single artists.
A different single-purpose visualization is Listening Clock[22] (see figure 4.2 (b)). This clock-shaped barchart is similar to the general purpose visualizations in its focus on the temporal dynamics of listening, but explicitly shows the listening intensity over the hours of a single day. The graph also makes a distinction between weekdays and weekends (cf. [121]) through color coding.
Several visualizations focus on last.fm as a social network. Last.Forward[23] creates a tree or a node-link diagram of friends and neighbors and allows filtering for some of the characteristics. The Last FM tasteOgraph[24] (see figure 4.2 (c)) displays a graph with artists and listeners as nodes that shows at a glance preferred artists that form the connection between two profiles.

All in all, both summarizing and single-purpose visualizations provide insight into listening histories, but suffer either from a too broad or too narrow approach: They answer a question either too vaguely or too specifically. By adding interactive filters or other adjustable aspects the listener is able to steer the exploration into one or the other direction, but only in a very restricted way (if at all). On the other hand, these W questions con-

[17] http://bandlogos.descentrecords.com/

[18] http://lastfm.gammalyrae.com/web/collages/albums

[19] http://stas.sh/lastfm/

[20] http://crosbow.peterpedia.org/lastfm

[21] http://playground.last.fm/demo/genderplot

[22] http://playground.last.fm/demo/clock

[23] http://lastforward.sourceforge.net

[24] http://geniol.publishpath.com/last-fm-tasteograph

densed into predefined visualizations keep the entry hurdles low and provide instant gratification directly after entering one's last.fm user name. Spreading these exploration tasks over an ecosystem of simple visualizations is certainly better than providing a complex all-purpose tool. However, the idea of instant gratification and the subsequent introduction of simple-to-learn tools for exploration can also be used to entice listeners into more sophisticated exploration. The examples in chapters 5 and 6 show how to approach casual visualization this way.

4.2 The design space of listening history data

In the following, I will discuss the important data dimensions of the design space of listening history visualizations. This design space is, of course, just a subset of the lifelog visualization and general information visualization design spaces and emphasizes the important data dimensions. There are other relevant aspects that should be taken into account, but are not explicitly addressed here, such the visual encoding of information (cf. [28, 173]) that determines how the available data is brought into visual form, or the role of interaction in exploration and sensemaking (cf., [88]). Another important aspect are the various tasks and goals that should be supported by the visualization system (from Shneiderman's low-level "overview first, zoom and filter, then details on demand" [151], via Zhou and Feiner's Visual Task Taxonomy [179] to Amar and Stasko's Knowledge Tasks [3]) and are versatile even in such a narrowly-defined space as listening histories. Finally, there is also the question of having an overarching process from design to validation (cf. also [3] or Munzner's more recent Nested Model [112]) and how to address the problem of evaluating infovis systems [131] (which is also the topic of active research, cf. proceedings of the BELIV workshop).
Solutions to these various aspects are presented in the respective discussions about the visualization prototypes. In this section, I will focus on the main data dimensions of *time*, *items*, and *listeners* and additional aspects (*background knowledge*, *device*, *task*) that determine presentation details.

4.2.1 Main dimensions

The design space for listening history data consists of the three dimensions *time*, *items*, and *listeners* (see figure 4.3). They determine how much data is taken into account and how much detail can be derived from the visualization. So, a hypothetical visualization that had *time* at a second/minutes granularity, *items* on the song-level and data from an individual *listener* would allow flipping through all minutes of the given data and displayed which song had been played at that time by the listener. It would, however, provide no function for summarizing this data by showing larger timeframes or abstracting from the song-level.

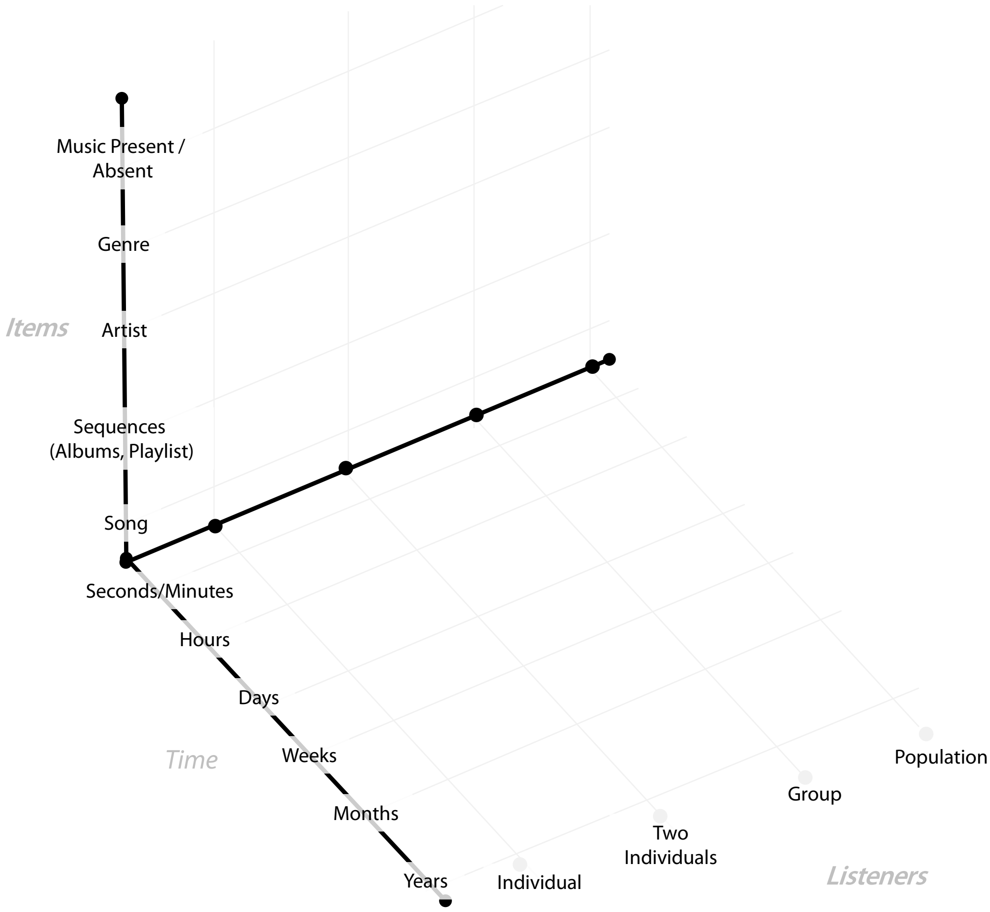

Figure 4.3 The main dimensions of the design space of listening history data: *time*, *items*, and *listeners*.

A visualization's position in the design space thus determines the level of granularity that can be reached but also which functions for summarizing or abstracting the data are available. A "perfect" visualization spanning the whole of the design space would thus allow accessing all listening data from an individual to the whole population with a granularity from single songs and seconds to genres and the presence of music and years. Of course, sophisticated interaction and visualization techniques would have to be used to make this amount of data graspable. Most visualizations focus on one dimension of the data.
What is not taken into account is what form of aggregation has been applied: It makes a difference whether song appearances are simply counted (as in the Last.fm Charts) or if each of them can be accessed and connected to a date. However, in the case of the charts the underlying timeframe can be adjusted, so access to the charts of the current week becomes possible, if only in a summarized form.
These main dimensions are based on the available real-world listening data (see section 3.3) and existing visualizations. As lifelogging capabilities and devices mature,

more and more of the ideal listening history data will become available allowing more sophisticated visualizations and emphasizing other dimensions (both task and social surroundings from the local context (see section 2.4.2) would constitute important dimensions, but are currently just not available). Also, these dimensions are based on the content of the data, in contrast to additional aspects that determine the presentation (see below).

Time

As music listening histories are a part of a person's life story, time is a central dimension for accessing and understanding the data. Integrating other data sources for hints of contextual information and supporting episodic memory also relies on time as an explicit dimension: Listening decisions become easier to understand when the temporal context is available (e.g., a stretch of stressful work or the holiday season) and longer periods of time provide a powerful way to aggregate data and find overarching patterns up to the development of musical taste over one's whole lifetime.
The temporal aspect can be presented with different levels of detail: They span from single *seconds/minutes*, over *hours* and *days*, to *weeks*, *months*, and *years*. Larger timeframes are imaginable, but we are currently lacking the necessary data to create them.
As with all dimensions, higher detail means being able to see more finely-grained patterns and finding explanations for very subtle occurences, but comes at the price of complexity. The temporal level of detail is also independent from that of the *item* dimension: Even if time is only represented in chunks of weeks or months, items can still be made accessible as sequences of single songs within that week or month (even though it might be hard to think of a suitable use case for this combination).
As both dimensions are independent, the time dimension can show repeating patterns, the overall frequency and intensity of listening and the length of listening sessions and single songs.

Items

Musical items or songs form the main part of listening histories. Each song item represents a specific memory, idea or impression for a person and is connected to various events in their lives. The *Item* dimension is based on the *relational* nature of songs. Their relationship can be either inherent (content-based similarity, genre- or tag-based similarity, created by the same or related artists) or develop in the "post-artifact" phase [111], creating the connection between two songs through an external event such as listening to them in succession (important for collaborative filtering and also for all prototypes in chapter 6), having them at the start of two different (important) playlists, etc. The important difference between the inherent and event-based relationships of songs is that the former is the same for all listeners, while the latter represents a very personal aspect and has to be determined and stored for each listener on their own. Especially the event-based relationship is important for understanding listening histories, and can

	Single songs	Song sequences (albums, playlists)	Genres and artists	Music present / absent
Duration	Long/short songs	Long/short sequences	Tendency to listen to the same genre/artist in a row	Length of listening sessions, noise/data gaps
Frequency	Intensity, most/least favorite songs	Intensity, most/least favorite sequences	Intensity, most/least favorite genres/artists	Intensity of music consumption, special events (parties, holidays), loss/gain of interest in music
Periodicity	Turnover rate for songs	Turnover rate for sequences	Flexibility/changes in taste	Daily/weekly/monthly routine (sleep-wake rhythm), changes in timezone

Table 4.1 Temporal patterns in listening history data divided into duration, frequency and periodicity (adapted from: [13])

be partially extracted from the existing data (e.g., how often a certain song has been heard in the last time and a corresponding estimate of its position within the song lifecycle). Other aspects such as the emotional attachment or the mental connection betweens songs and external events (e.g., vacation hit, song of a personal relationship), however, cannot be easily derived from real-world listening data and have to be gleaned from a listener's memory.
The *item* dimension can also be used to reduce the complexity of a listening history. Similar to the level of detail of *time*, a corresponding scale for items is the musical hierarchy (see section 3.3.2 and [13]). Possible levels of detail are thereby *single songs*, *musical sequences* such as albums and other predefined playlists, *artists*, *genres* and finally, at its most abstract, the *presence/absence* of music. These levels of details can again be used to reduce complexity or intensify insight.
In combination, the *time* and *items* dimensions can be used to find recurring patterns. Depending on the level of detail of both, different insights are enabled. One way to classify patterns in temporal data was proposed by Laxman and Pastry [95], who divide events into *duration*, *frequency*, and *periodicity*. Using this classification, we can identify temporal patterns in listening histories and organize them based on the level of detail of the item-dimension (cf. [13]). An overview of these patterns is shown in table 4.1.

Listeners

As we have seen in section 2.1.3, identity construction is one of the central uses of music. Self- and other-directed forms of identity management (cf. [56]) are important for music and accordingly, drawing conclusions from people's taste is a popular pastime. But listening histories not only contain information about items and musical taste, but also about time and thus listening behavior, and therefore allow comparisons regarding the intensity of music use or of recurring patterns (e.g., whether two people listen to music on their daily commute).

The *listener* dimension again has several levels of detail that determine how much aggregation happens with the underlying data. On the lowest level, an *individual* listening history is shown. *Two individual* histories are also common for direct comparison, while less detailed cases contain either *groups* of histories or even span large parts of or the entire *population*.

Comparing one's own listening history with others can help in finding common or disparate taste, coming up with topics for conversation and discovering new music. Therefore, most visualizations either have a single listening history or two histories side-by-side at their core. Adding one more history to an existing chart (e.g., a timeline) still allows performing all intended actions without requiring additional filters or other techniques. Everything beyond two histories, however, can make things much more difficult and trying to display the listening behavior of a whole population in maximum detail is practically impossible. The listener dimension can, however, be simplified by aggregating listeners based on their social relationships or demographics.

The three main dimensions *time*, *items*, and *listeners* provide a classification for visualizations and uncover their strengths and weaknesses, just as the emphasis they put on certain aspects of the data. Figure 4.4 gives an overview of the visualizations from section 4.1.1 in the context of this design space. It is clearly visible that all of them rely on larger timescales and more abstract music classifications such as artists and genres, i.e., low levels of details on the time and item dimensions (with the exception of the Last.fm Charts, that are able to provide a list of the most popular songs of the last 7 days). This lack of more detailed information makes the summarization easier understandable and the visualization clearer. It also requires less interaction. It is interesting to note that these visualizations either have a single or two individual histories or data for the whole population, but do not merge them by, for example, providing popular opinion as a backdrop for one's own listening. Single-purpose visualizations (section 4.1.2) operate on similar levels of detail for time and items but sometimes provide this backdrop: They use the number of listeners for an artist (Underground-O-Meter) or their demographics (Gender Plot) to "measure" the position of an individual listener within the population. Depending on the task they are supporting they might also rely on different ways to aggregate songs (e.g., distinguishing between weekends and weekdays and counting songs per hour in Listening Clock), ignore timeframes (e.g., just taking presence or absence of artists for finding common taste in tasteOgraph) or completely shun visualizing the listening history (e.g., by interpreting last.fm as a social network in Last.Forward).

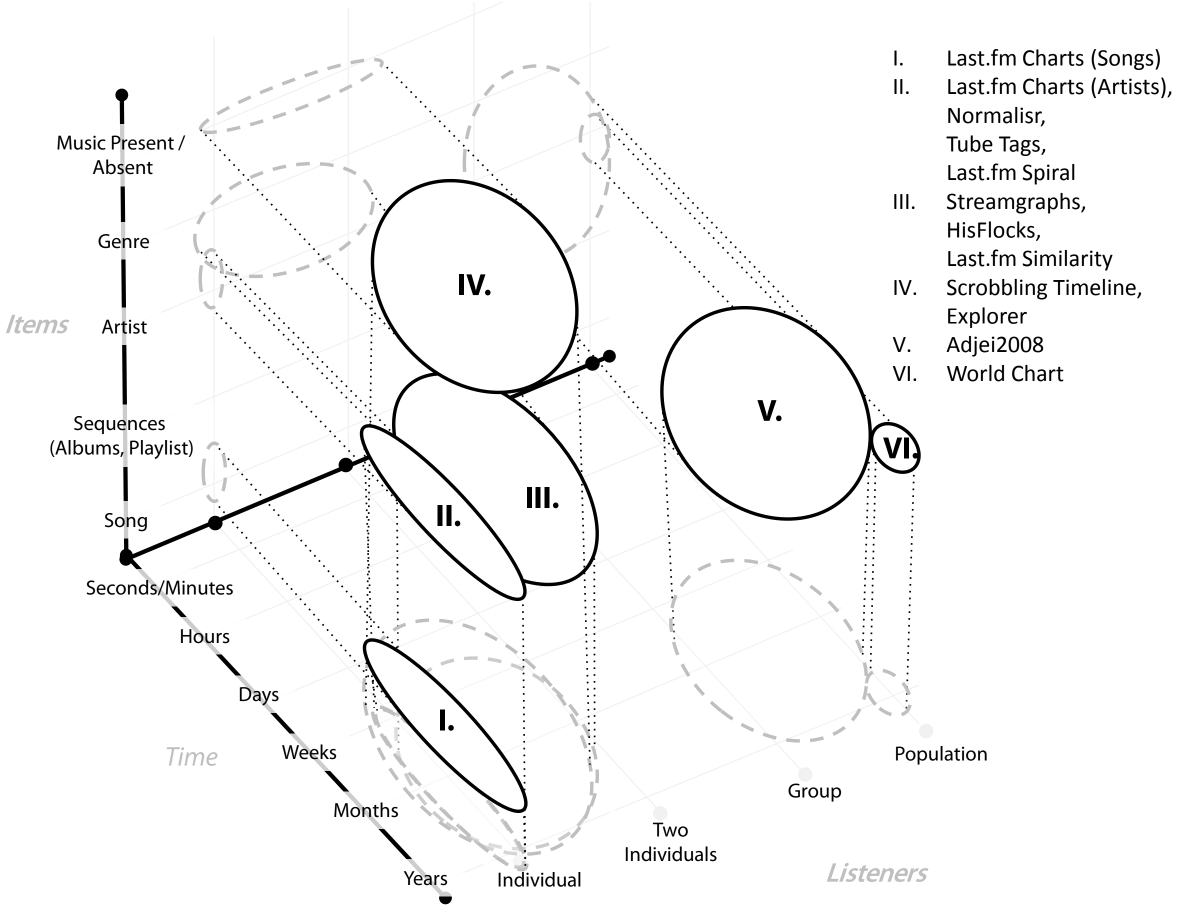

Figure 4.4 Summarizing visualizations in the design space of listening history data. Their focus on aggregating the time and items dimensions is clearly visible.

4.2.2 Additional aspects

While the main dimensions provide a way to categorize listening history visualizations, they are not the only influence on the resulting tools. In the following, I will describe three different aspects that are also important for the way the information is presented. These are the analyst's *background knowledge*, the *device* used for accessing the data, and finally the *purpose* of working with the data.

Background knowledge

An analyst's background knowledge determines what presentation is best for her or him in working with a listening history. One aspect of that is *musical training* that determines

how specific descriptions that are added to the visualization can become. While the musical genres that have been described above might be helpful for almost all listeners, concepts such as mode or key or the workings of chords and harmony might feel more arcane. Similarly, relying on background knowledge about certain artists in the description of a song or band ("Rooney ... sounds like the love child of the Cars and the Beach Boys"[25]) can be difficult to understand for music novices or non-cognoscenti of a certain subculture. On the other hand, keeping the level of background information as general as possible can also frustrate experts: If the purpose of a visualization is to support fans of classical music in exploring their listening histories, relying on genre descriptors and artists instead of periods and composers can be disappointing and, in the worst case, render the tool completely useless. Sometimes, other ways of abstraction or an addition of more complex or uncommon musical attributes can be beneficial.
A second aspect of background knowledge is whether the purpose of the visualization is to explore one's own listening history or that of another person, i.e., if *personal background* is available or not. Supporting episodic memory through suitable triggers can only be fruitful if this memory exists. On the other hand, if someone else is looking at an external history providing as much of the listener's background as possible can be extremely helpful, while it would be mostly annoying for the actual owner of the history. Scenarios where external histories are used are common for tools that either contrast or summarize multiple histories.
The last important type of background knowledge is *experience with visualization tools*. Similarly to the musical expert above, a visualization expert can be frustrated with a tool that does not support complex enough manipulations of the data or provide functionality that she or he knows to exist. Visualizations that generate static images and maybe even distort the data for aesthetic reasons (Tufte's "lie factors" and "chartjunk" [161]) might find fans with no statistical or other scientific background, but can be repulsive for others. Therefore, knowing the intended audience of the visualization and its music and visualization experiences is central to creating the right visualization tools.

Device

Another necessary deliberation is that of the target device. With the proliferation of music listening to handheld devices, supporting tasks around listening histories on-the-go is important. Yet, the available screen dimensions, the mode of interaction and the possible surroundings make a large difference for the resulting tool.
Relying on a *smartphone* as a device for visualizing a listening history restricts the available screen size so that some of the successful desktop visualization techniques simply might not work (see, for example, [57]). Additionally, restrictions in battery performance, processing power and network access can make a difference for tools. Smartphones nowadays also use touchscreens for input, which require suitable adaptations: Fingers create occlusion and make it harder to hit small targets (cf. [168]), which means

[25] http://www.vh1.com/music/you_oughta_know/rooney/main.jhtml

that widgets containing relevant information have to be enlarged and pixel-perfect targetting is not necessarily possible. In addition, smartphones are used on-the-go, which means that their usage style is different from that of a regular desktop PC: Interruptions are the rule and either arise from the circumstances (getting off the bus, talking to somebody) or the device itself (incoming calls or messages, calendar reminders). Looking at the phone display while performing another main activity (e.g., driving a car, walking) or in short bursts of otherwise unused time (e.g., while commuting) means that smartphone tasks usually have a lower priority and are restricted in their duration. Therefore, single-purpose tools are usually more suitable than complex analysis applications.
The classical *computer*, either in portable or stationary form, provides a high-resolution screen and complex input through mouse and keyboard, which allow for equally complex visualization tools. Increased processing power and (usually) faster network connections make it easier to work with larger data sets. The usage style is also different, as using a computer requires a more conscious decision than playing with the phone to pass the time. Therefore, visualization designers can rely on more leeway for idle time and tasks that require more than a five minute session. The recent rise of web-based computing also made it possible to deploy applications directly over the web or even run them in the browser, which created new ways of distribution and receiving feedback (see chapters 5 and 6 for examples).
An interesting upcoming third form are *tablet computers*. Tablets in their current incarnations share many characteristics with smartphones but are different in other central aspects. They similarly have less processing power, depend on batteries and touch-based input with its resulting problems, but are different in their usage style: While tablets are portable and can be used on-the-go, they are also commonly used as "couch computers", in a more laid-back fashion. Compared to smartphones, this allows more complex tasks which is also supported by the larger screen sizes and higher resolutions.
As smartphones, desktop and laptop computers and tablets all have capable web-browsers that support novel standards such as HTML5, future visualization tools might work on all devices simultaneously. Yet, they should still take into account the differences in capabilities and usage style.

Purpose

A last aspect that determines the design of a visualization tool is its intended purpose. For a simple, static visualization that is geared towards *self-presentation* (e.g., HTML-embeddable badge images) interaction is much less necessary than bringing a point across (e.g., one's infatuation with Justin Bieber). Yet, with listening histories as lifelogging data, *analysis* or reflection tasks can benefit greatly from supporting more complex interactions and data mining.
Another purpose for using listening histories can be the support of *reminiscing*, to "relive past experiences for emotional or sentimental reasons"[149]. Such reminiscing tasks do not necessarily require visualization and might be supported in a much more subtle way: Cosley et al.'s *Pensieve* system [35], for example, sends snippets from listening

histories or other social media via email to prompt reminiscence. Similarly, supporting the *rediscovery* of forgotten songs can also happen in different contexts. In our own *SongSlope* (see section 6.3) rediscovery is supported within a media player instead of an explicit visualization tool.
Finally, the purpose of using listening histories can also be to support standard *music tasks* (see section 2.2), such as listening (e.g., for generating playlists), or organizing a collection (e.g., by deriving similarity from side-by-side listening occurences). Discovery can also be helped by deriving taste from histories.
All in all, the intended purpose is central for the capabilities of the system and also the form it takes.

4.3 Summary

In this chapter I have discussed the topic of listening history visualizations. I first gave an overview of existing visualizations, that were mostly created by last.fm themselves or enthusiasts. They can be divided into summarizing visualizations, that give an overview of listening activity or taste, and single-purpose visualizations, that support a single task efficiently. I then presented a design space for listening history visualizations that consists of the main dimensions *time*, *items*, and *listeners*. Depending on their position within this space, visualizations apply different forms of aggregation for the data and provide different levels of detail. Other aspects that are important when designing visualizations in this space are the analyst's *background knowledge*, the used *device*, and the intended *purpose* of the tool.
The next two chapters present examples from this design space that I created in the course of this thesis. Chapter 5 emphasizes the *lifelogging* aspect of listening histories and presents visualizations for analysis and reminiscing. Chapter 6 focuses on the *relational* nature of songs and suitable tasks such as building playlists and rediscovering forgotten favorites.

Chapter 5

Listening *histories*:

Visualizations for analysis and reminiscing

Such longing for the past
for such completion
What was once golden has
now turned a shade of grey.

– Bloc Party - *Two More Years* –

Listening histories are purely vanity data, compared to some other lifelog information. They are certainly nice to have and can bring interesting insights into one's taste in music and relationship with other people, but not having them does not prevent important activities or impede them in any way. Other lifelog data such as medical readings and logs can help in losing weight or becoming healthier, but listening histories are, just as status updates in social networks or check-ins in Foursquare, nothing more than nice to have.

Still, last.fm's popularity and millions of listening profiles show that there must be something that keeps people interested and investing into logging their music listening. Similarly, the lifelog movement that explores all areas of human experience and not only the "useful" ones of health services, attracts a large following that captures meals eaten, clothes worn, TV shows seen, and distances from home to work and back. When looking at their antics, one cannot help but wonder: What for?

Beyond the neat tasks supported by lifelog data that apologists like to bring forward when asked about their motivations, there is one central underlying idea driving all of lifelogging: Narcissism (hopefully mostly in the healthy sense). The nature of man as a conscious being necessitates an interest in the self. Denying this interest in one's own existence is either detrimental to one's general motivation and can lead to completely apathy, or simply false: Everybody likes to hear stories about themselves, even if modesty keeps us from asking for them.

Also, more knowledge about one's character traits and flaws can help us become better people. It is no coincidence that the Greek aphorism "Know Thyself" was inscribed in the mystical oracle's temple at Delphi [105]. As in all of life, more knowledge allows making better and more informed decisions and reaching one's goals easier. Understanding how to support flow [36] in one's work, what depressing songs to avoid and what job to pick to achieve happiness might be hard (cf. [144]), but certainly becomes easier with enough background knowledge. Indeed, all of music's mood management functions require a solid understanding of its impact on one's emotions.

So, listening histories are still vanity information, but this vanity is only human. And sometimes they might even be helpful for something else.

In this chapter, I present four visualizations for analyzing and reminiscing. These purposes distinguish between a more rational approach to working with listening histories and a more emotional one (cf. [13]). Analysis is the more rational use case, that "describe the typical tasks of finding patterns and testing hypotheses based on a set of data" [13]. As these patterns should be inherent in the data, an analysis can also be performed by an external observer and not only the owner of the data. Memory triggers can be helpful but are not essential.

Reminiscing on the other hand cannot be performed by somebody who is not as intricately connected with the depicted songs as the creator of the listening history. In this case, songs work as memory triggers that allow a person to relive parts of their lives and a listening history is more like a photo album that holds dear memories encoded in song.

The visualizations in this chapter move between these two poles and address the problem of making a listening history understandable. They all allow access to single songs, but provide mechanisms for abstraction to make sense of bigger time periods. Even though they do not emphasize the *item* dimension (see next chapter), they still contain hints at the connections between songs. While they mostly focus on understanding they can also be used to actually listen to music. And they all contain ways to interact with the data and actively explore it, but provide immediate benefits directly after launch.

These aspects and similarities are visible when shown in the design space: Figure 5.1 shows all visualizations of this chapter organized along the *time, items,* and *listeners* dimensions. Their *time*-centric nature provides a high level of detail for this dimension and insight about the presence/absence of music for all of them. Yet, they also show (mostly) detailed information about items (from songs to genres). They are, however, divided along the *listeners* dimension, as two of them only work for individuals.

In the following, I will describe the individual design processes and results for each of the visualizations and also evaluations (when appropriate).

Strings is a playful visualization for an individual listening history that enables focusing on interesting items through interaction.

LastHistory is more powerful, but hides its complexity through a reduced and aesthetic interface. It explicitly separates usage modes into analysis and reminiscing and to this end, integrates photos and calendar information as memory triggers for episodic memory.

LoomFM emphasizes the relationships between two individual listening histories, also in a playful and minimal way. The level of detail about songs is coupled with the visible region of the history.

LastLoop again provides powerful analysis capabilities for understanding histories and supports a variable number of them. Therefore, it allows exploring an individual listening history or finding similarities in taste between multiple ones.

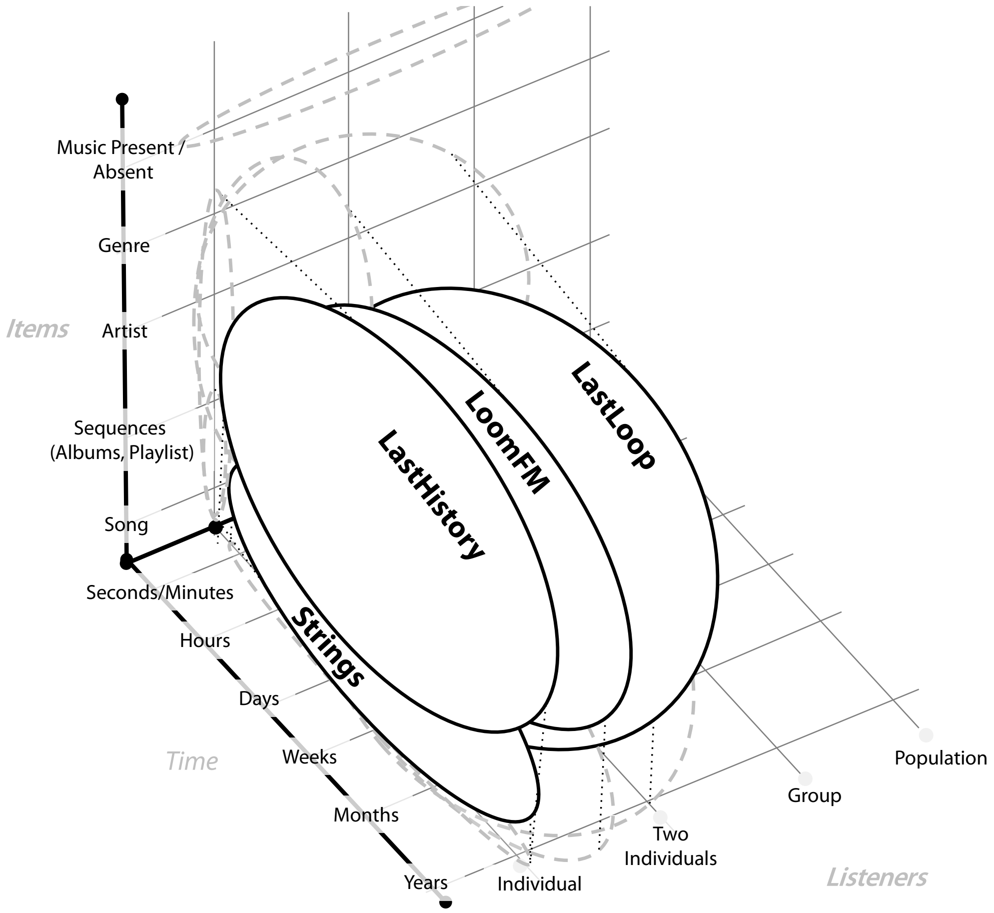

Figure 5.1 The four visualizations presented in this chapter in the listening history design space. All four are able to present longer periods of time, different levels of abstraction and support individual or multiple listeners. Note that all four also support insights on the presence or absence of music.

5.1 Strings

The first questions that owners of a listening history might ask themselves when trying to understand their behavior are very basic in nature: How much music have I heard? When exactly? For how long at a time?
All these time-centric questions are only insufficiently answered by last.fm's default view for listening histories. These long, table-based lists of songs (see figure 5.2) are helpful for retrieving the last tracks heard or looking into specific timeframes, but provide no overarching understanding or enhance finding interesting outliers. The textual timestamps make it difficult to estimate the length of listening sessions, or see long stretches of time where no music was heard.

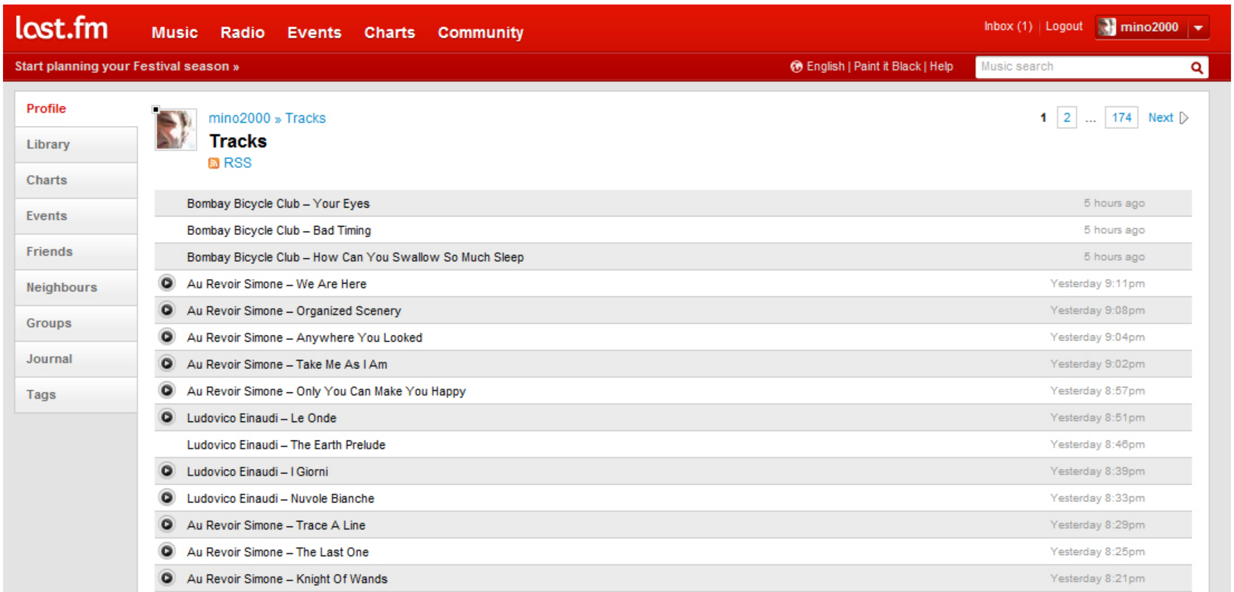

Figure 5.2 Last.fm presents one's listening behavior as long lists of songs (source: last.fm)

I thought of these time-centric queries and understanding when designing *Strings* (see figure 5.3 and [12]), a simple visualization that provides insights into these questions.

5.1.1 Population and goals

One of the additional aspects of visualization design (see section 4.2.2) emphasized the *background knowledge* of prospective analysts and the purpose of their analysis. In user-centered design lingo (see, e.g., [34]), these constitute the user needs and goals.
At the beginning, my goal was to create a replacement for last.fm's lists. Therefore, the average "analyst" using the tool would be the owner of the displayed listening history her- or himself and no other histories would be included into the visualization. As

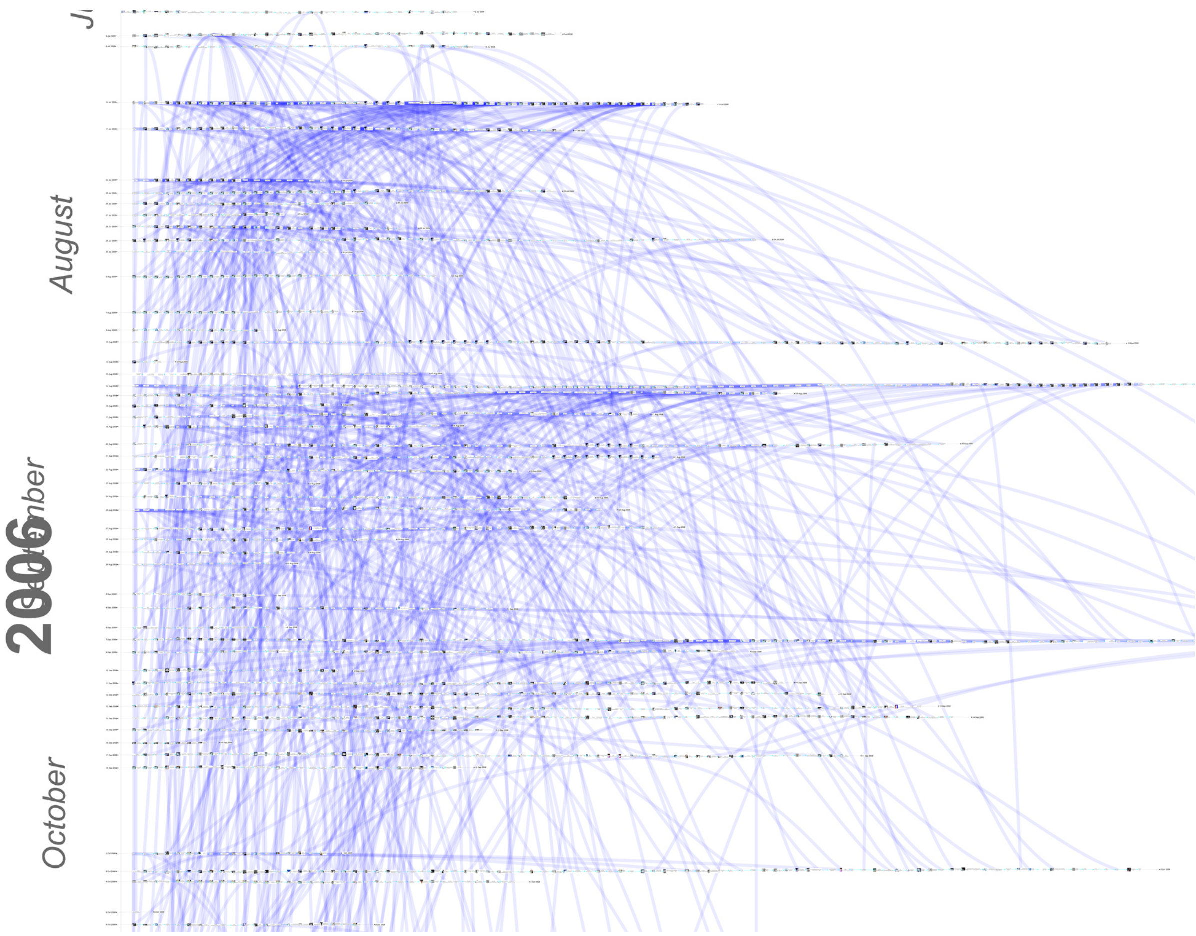

Figure 5.3 *Strings* is a playful and time-centric visualization that emphasizes listening sessions.

last.fm's population is vast, only vaguely defined and certainly does not fall into musical stereotypes, I could not expect a solid background in music. Most importantly, I tried to restrict all necessary knowledge to that already contained within the data, i.e., time stamps and song/artist names. Similarly, no background knowledge about visualization techniques could be assumed, so I would try to keep the complexity of the interface as low as possible. Seen from the perspective of the analyst, she or he has the following goals:

A: Find temporal patterns on the song-level (implying a focus on the *time* dimension of the design space). These patterns work without information about the underlying songs and only rely on information about the presence or absence of music: As described above (see right column of table 4.1), they include the general amount of music, gaps in the data, the frequency of listening and its change over time, and daily routines. The relationship between songs and time thus had to be clearly visible.

B: Find patterns along the *item* dimension. While the main focus would be on time, relationships between time and items or among items should also be made visible as well as possible (without interfering with the main goal). Especially repetitions of single items were of interest.

C: Have easy-to-learn interaction and an immediate benefit from the visualization. As the analysis of one's listening history would be something that was not performed regularly and at best repeated every few weeks, keeping the entry hurdles low and providing an immediate benefit would be important. Also, the available screen real estate should be used for the visualization and not the interface.

Based on these main goals I designed the *Strings* visualization. The main goal of uncovering temporal patterns in a single listening history defined the central focus of the visualization.

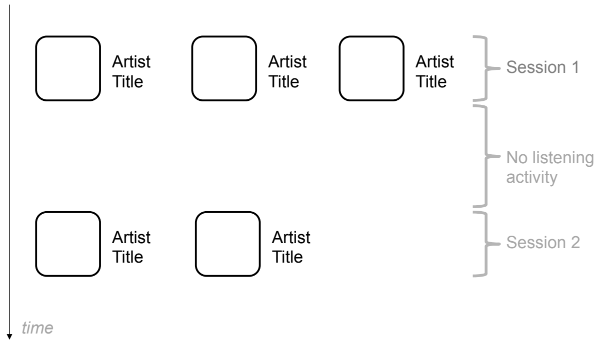

Figure 5.4 In Strings, listening sessions are arranged along the timeline, providing insight into session-based patterns.

5.1.2 Design

Timelines have proven to be beneficial for uncovering temporal patterns in personal histories (e.g., [132] or [166]). Therefore, I decided to use them in *Strings* and arrange all songs along a timeline. This timeline would be vertical, similar to last.fm's table-based view with the oldest songs at the top. As the timeline would span even on average several months and fill multiple screens, simply displaying it would not be enough to give an overview. I used a zoomable user interface (cf. [14]) that would allow dynamically focusing on certain regions of the timeline and getting an overview. In addition, a ZUI would keep the footprint of the interface low (goal C).

The next question was the granularity of the timeline: Songs only span several minutes and to show them in detail (goal A), the scale of the timeline had to be similar. This, however, would have become a problem in connection with the overall length of a listening history: Displaying multiple months in one-minute-granularity would have resulted in

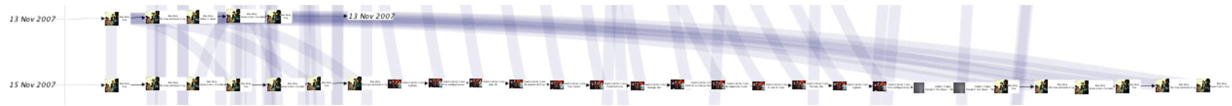

Figure 5.5 Other instances of a song are connected via arcs which can uncover repeating playlists.

a very long (and empty) timeline. Deciding on an arbitrary cut-off (e.g., one hour), on the other hand, would have been of only limited use, as it would have separated sequences of songs that started right before the end of the hour etc. Listening sessions are suitable separators for listening activity (cf. session 3.4), so I used them in *Strings* (see figure 5.4): All songs in sequence that are separated from the next song by less than one hour belong to the same listening session. The sequences of songs within one session would be displayed horizontally, thus allowing comparing their lengths and seeing patterns simply by looking at the horizontal dimensions. Also, the formerly empty horizontal space would be (potentially) filled with information. Stretches of time void of listening sessions would be shown but not true to scale. This allows seeing empty parts of the history, without taking up too much space.

In order to keep the required background knowledge as low as possible, single songs just showed the artist name and title. Additionally, I used pictures of the artists from last.fm as visual indicators for recognizing songs at a glance.

One main problem of using such a time-centric logic for the layout are multiple instances of single songs. As every song instance appears on its own somewhere within the history, their visible representations are also scattered throughout the timeline. While *Strings* was centered around the temporal aspect on purpose (goal A), I still wanted to add some information about song appearances and the items-aspect (goal B). Therefore, the first instance of each song is connected to all its other instances with blue arcs (similar to arc diagrams [175]) (see figure 5.3). This does not allow seeing all other instances at a glance (only if the history contains very little repetition) and can degenerate into an indiscernible tangle for multiple repeating instances, but it provides hints about this information without interfering with the time-based layout. It also works very well for seeing certain patterns: Spatially close sessions with multiple repetitions become entangled in arcs, and repeating playlists are clearly visible (see figure 5.5).

In order to navigate the timeline some interaction is necessary, even though minimizing it was a main goal (goal C). All interaction happens with the mouse: Panning the timeline is possible through dragging the mouse and zooming in and out happens with the mouse wheel. This is enough to navigate the information space and switching between receiving an overview or focusing on a certain region.

As the arcs that connect instances of one song are not meant to and mostly do not work for visually following them, I decided to rely on interaction for this use case. Each song is part of a series of sequences (see figure 5.6). By clicking on a song, all other instances of it throughout the timeline fall into the same position, taking the respective sequences with them. While the songs' positions are usually static, their layout is then determined by a force-directed positioning algorithm which avoids overlap.

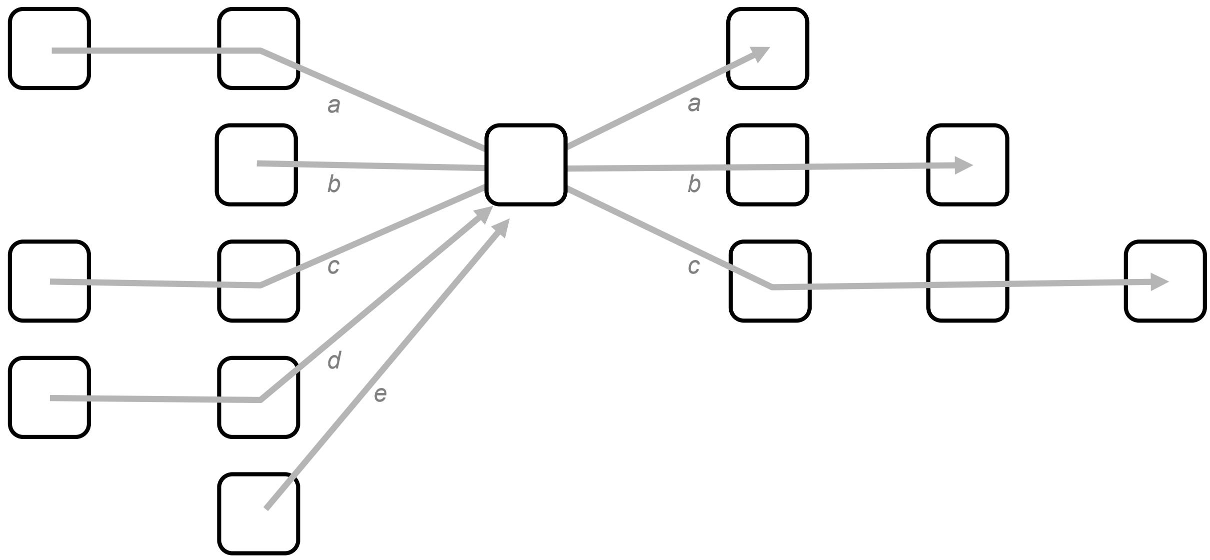

Figure 5.6 By clicking on a song in *Strings*, all sessions where it appears are shown (in this case labelled *a* to *e*), making its predecessors and successors immediately visible.

5.1.3 Discussion

Strings was a first foray into listening history visualization and originally intended as a simple to use replacement for tabular representations. One of the goals was to reduce necessary interaction and the visualization does that by providing an immediate benefit: Without interaction and directly after launch, the songs' positions on the timeline already exhibit multiple temporal patterns: Empty stretches of time become visible, just like up- and downswings of listening activity and the length of listening sessions. The underlying blue arcs that connect instances of songs show more intense areas of repetition and playlists. Navigating the timeline using zoom and pan is trivial and clicking on songs shows containing sequences.
Strings was implemented using Heer et al.'s *prefuse* framework (an InfoVis programming framework for Java)[67] that provided capabilities for managing the data and several layout algorithms. While powerful for predefined functions and layouts, prefuse proved to be complex to adapt for different use cases and shaping the results of the layout algorithms took some time. Also, Java's performance allowed only up to around 2,000 songs while keeping the interaction fluid.

5.2 LastHistory

Strings had shown that uncovering temporal patterns within listening histories was easy to do using a timeline. This came, however, with a strict focus on time, as the connections between song instances where mostly of only aesthetic use. With our [1] next project, *LastHistory* (cf. [13] and see figure 5.7), we planned to have a stronger focus on more complex interaction capabilities and more support for exploration.
An aspect that we felt had been underrepresented in *Strings* was support for understanding the timeline. While dates might be helpful in roughly mapping a song event to a time in one's life, we wanted to make sure that listeners had enough reference points to immediately understand what they were looking it. For this, we especially thought about how to integrate suitable memory triggers (see section 3.5.2).
Another shortcoming of *Strings* was the missing formal evaluation. While we had received informal comments, the lack of a coherent frame for feedback kept them very superficial. Therefore, we also wanted to have a suitable validation phase for *LastHistory* to see if our theories about immediate benefits and non-threatening interfaces were valuable. And we were also interested if there was an actual need for analysing one's listening history.

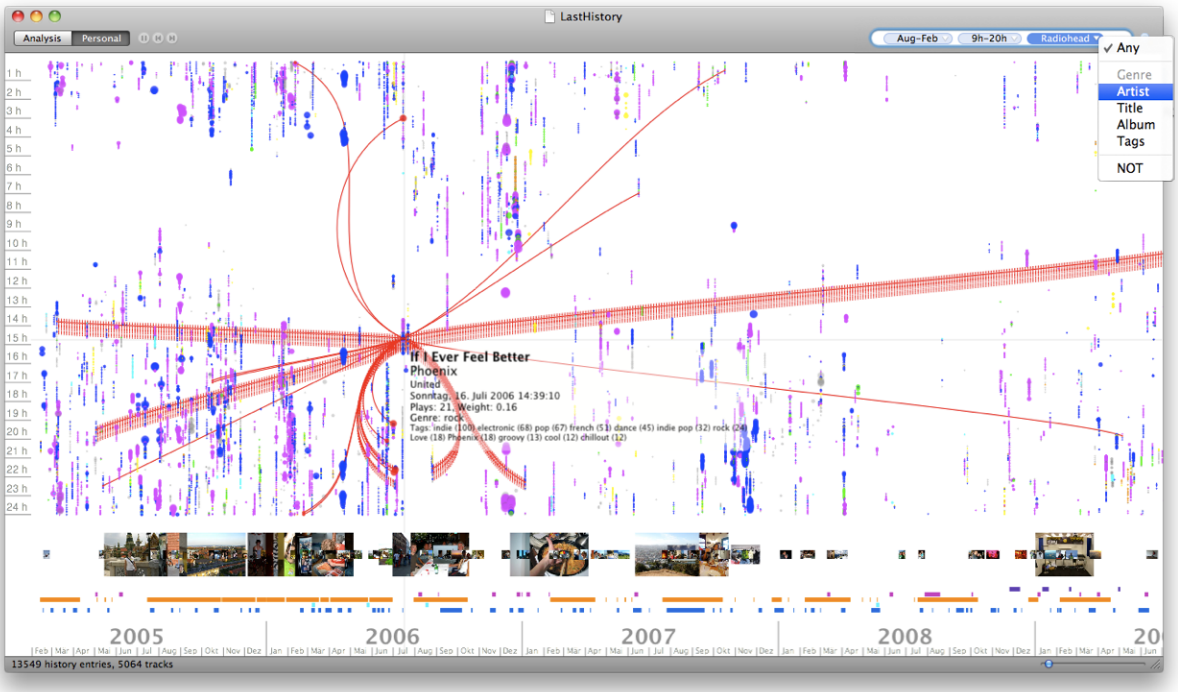

Figure 5.7 *LastHistory* is a time-centric casual visualization tool for analysis of and reminiscing with music listening histories (Source: [13]).

[1] *LastHistory* was part of Frederik Seiffert's Diplom thesis [147].

5.2.1 Population and goals

LastHistory was planned from the beginning as a casual visualization, which meant that the evaluation also required real-world, casual participants. Therefore, the population that the tool was aimed at was again very diverse. Basically everybody who owned a last.fm account should be able to use it, so we could not expect much background knowledge in either music or information visualization. We would, however, try to support such knowledge if present. From early on, we also planned to make the tool available online and try to get feedback. So keeping it simple was paramount, as we certainly did not have the resources to provide the necessary support for a complex application.
Regarding the user population, we made one decision that was different from *Strings*: While the latter had only been aimed at the owners of listening histories who would look at their own data, with *LastHistory* we also wanted to help external analysts. At the time of development, we were ourselves looking at various listening histories trying to understand the types of relevant patterns within them. Unfortunately, tool support for doing that was scarce - visualization tools were either too general-purpose (e.g., Tableau) or too specific (see section 4.1 about other visualizations). We wanted to have a tool that allowed drawing complex conclusions, but was convenient too use and easy to feed with data. Such a tool would also be helpful for other researchers (e.g., musicologists or ethnographers) once they started working with such data.
Overall, however, the goals of both groups were mostly similar:

A: Find temporal patterns on the song-level. Supporting any kind of reminiscing requires binding the data to time, which is why decided to make temporal patterns (again) the main goal of the visualization. In this project, we wanted to support the full spectrum of temporal patterns, which only works in combination with the item dimension (see table 4.1).

B: Have support for exploring the *item* dimension on the song-level. This was a necessity for the whole spectrum of temporal patterns. The focus again lay on repetitions of single items but also on the repetition of sequences (playlists, albums). Also, finding out about specific aspects of certain groups of songs (e.g., by one artist or of a certain genre) should be supported.

C: Use the listening history for analysis or reminiscing (or both). Depending on the intentions and whose listening history is displayed (own or external) the goal of using the tool would be either to look for patterns and try to understand what happened, or to reminisce and remember a stretch of one's life. Of course, for one's own history a mixture of these two sets of tasks would also be possible. These use cases should neither interfere nor complicate the interface.

D: Listen to music. Apart from the two use cases, a listening history is still mostly about music. Therefore, actually listening to it and using the tool for rediscovery of single songs or whole sequences would be desirable for both analysts and reminiscers. We had also repeatedly heard this suggestion for *Strings*.

E: Be able to rely on one's background knowledge. Although it is not a given that people are proficient in either music or visualization, if they are they are happy to make use of it. Even worse, if they know that some kind of information or interaction is available or possible, they can become frustrated by the lack of it and just abandon the tool. As mentioned above, interacting with one's listening history is nice to have but not essential, so the amount of frustration from the interface tolerated will be rather low. Another piece of background knowledge that becomes important for the reminiscing use case are memories and should also be supported if available.

The goals relevant to *LastHistory* incorporate the ones from *Strings* but extend them further. A main distinction is between people who look at their own, versus people who look at an external listening history. While the analysis goals are similar, reminiscing, relying on memories and listening to music will be possible or desirable only for one's own history. This led us to making an explicit distinction between these two approaches. We called them *analysis* and *personal* modes and adjusted the interfaces accordingly (but consistently) to hide unnecessary functions and make better use of the available screen space. Some functionality, however, is available in both modes, even though it might be more useful to the personal use case (e.g., song playback is possible in both).

5.2.2 Design

The order of goals suggested a clear priority for *time* over *items*. We used a timeline similar to the one in *Strings*, but aligned it horizontally instead of vertically. Instead of concentrating on listening sessions, however, we wanted to take repetitive cycles more into account. Based on hints from the literature (e.g., [120]) and our own experience with listening histories, we decided to emphasize daily cycles. Regular activities such as commuting, working, spending the evening at home should become visible through a suitable display of items.
The timeline is therefore not one-, but two-dimensional: The x-axis lists the days, while the y-axis shows the hours and minutes of the day when a song was played (cf. [13]). Similar music listening that happened on each day at around the same time (e.g., commuting) would be clearly visible, just as the distinction between weekdays and weekends (no music). The alignment to hours brought another benefit because of the way last.fm treats time. Internally, all time stamps are kept in UTS format (see section 3.3.2) and thus within the same time zone. In *LastHistory*, time stamps are used unmodified, which means that all times are shown in the UTC±0 time zone. Therefore, listening histories from different time zones have different patterns, even though overall daily listening is relatively similar worldwide (see section 3.4.2). Additionally, when changes

in time zones happened within a single history, they would also be immediately visible. Half-year internships on another continent usually also left their marks within the histories. Several examples for the visibility of daily patterns and the influence of the time zones are in figure 5.8 that show especially well sleep patterns (top row) and different time zones and noise in the data (bottom row). This approach to the song layout also

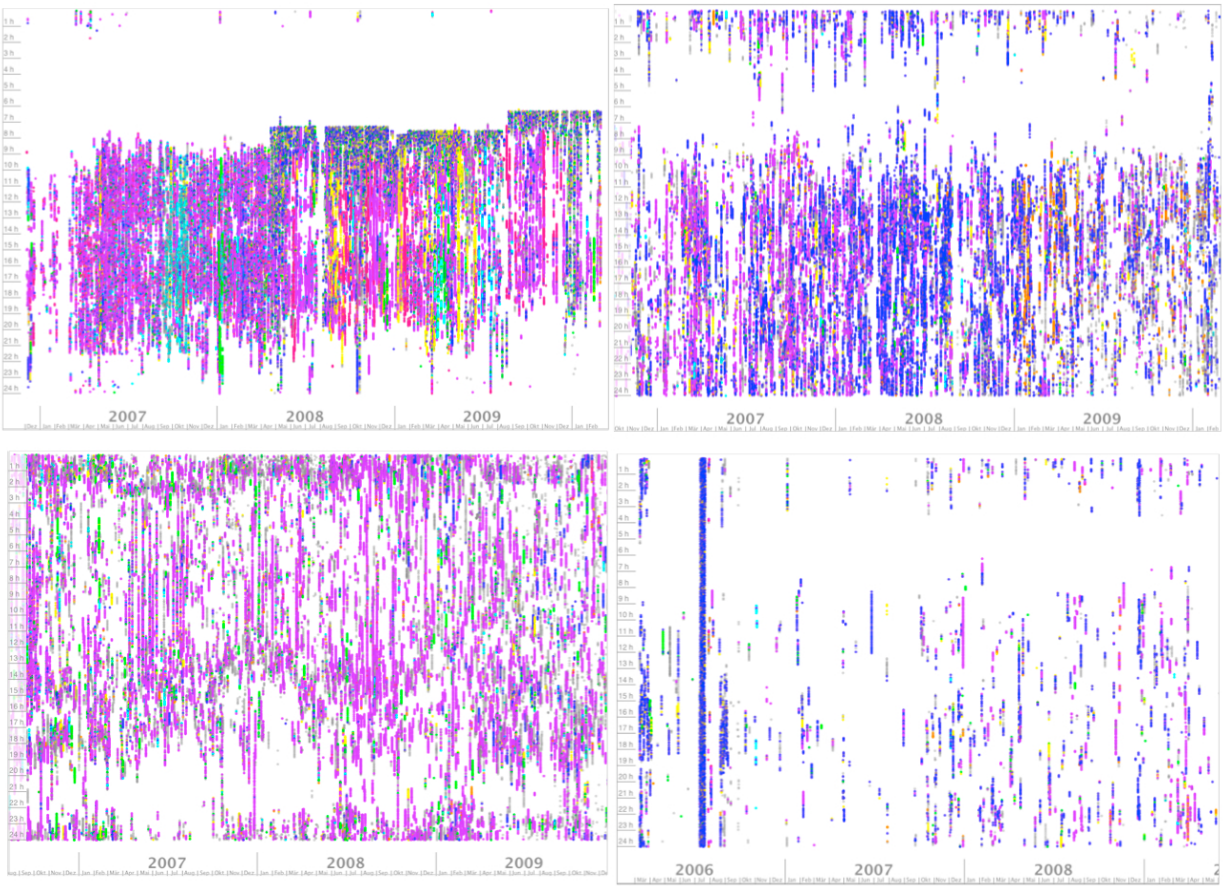

Figure 5.8 *LastHistory*'s two-dimensional timeline shows minutes and days which uncovers daily patterns such as commuting and sleeping (top row), different time zones (bottom row left) and even noise (bottom row right). (Source: [13]).

provided us with an immediate benefit as the visualization itself uncovered these patterns without requiring interaction. Other uses for the y-axis would have been possible (e.g., "other frequencies, playcounts, genres, tags, etc." [13]), but the secondary temporal dimension promised to trigger the most aha-effects without asking for any settings beforehand. A shortcoming of the fixed time distance of days was the cut-off that could split a listening session that started before 12AM UTC into two. The ease of discovering daily patterns made up for this downside of the absolute layout, however.
Several considerations went into the representation of songs. Similar to *Strings* and in order to fulfill goal B we wanted to show single items instead of an aggregated version of them (as, for example, in Streamgraphs). In *LastHistory*, however, we were very

restricted in the screen space available for displaying songs. As the y-axis showed 24 hours of songs and should be completely visible on-screen without scrolling, we had to make sure that the experience would not be ruined by too much overlap. If in theory all 24 hours of a day would be filled with music and a song had an average length of 3.5 minutes that would result in around 400 items. Song icons had therefore to be small and a size of several pixels at most (*Strings'* song icons took up around 60 pixels vertically for the artist image and text). We decided that some overlap would be in order, as the extreme case of a day filled with non-stop music would be rare. To make the best use of the available space and reduce visual clutter, our final song icons were simple shaded circles with a diameter of 5 pixels (the screen real estate of the y-axis depended on the window size: widget and status bars needed around 100 pixels, which left 800 pixels for the songs on a 1440x900 display). These circular song icons also did not convey the duration of songs, but we decided to sacrifice that for less restlessness, possible errors in the actual duration of listening, and the loading time of acquiring data about song durations (cf. [13]).

Songs encode their genre using color coding. As an accurate distinction between multiple colors can be difficult (cf. [173]), we tried to keep the number of different colors as low as possible and their hues far apart across the HSV spectrum. In the end we had eight genres and corresponding colors: "orange (classical), yellow (jazz), green (funk), turquoise (hip-hop), blue (electronic), purple (rock), pink (metal), grey (unknown/other)" [13]. A song's genre was determined based on the most popular of last.fm's user-generated keywords and for each genre we created a list of suitable tags ("'funk', for example, encompasses 'soul', 'disco', 'ska', 'reggae', 'worldmusic', etc" [ibid]). While seven genres seems a small number, the lack of a consensus regarding genres and their definitions (cf. [7]) meant that discussing about sub genres and which song belongs to which genre would have caused more trouble than brought insight. Also, we used a static mapping that had the same colors for every listening history, even though that meant that a rock-fan mostly saw only purple circles which questions the use of the color-coding. The decision to use a fixed genre-color-mapping came from two reasons: First, all visualizations of histories stayed comparable and showed at a glance which genres are the most important ones. And second: Once a person had learned the simple mapping between seven colors and genres they could rely on that without having to look up the meaning every time anew.

For the interaction, we wanted to keep *LastHistory* powerful but easy-to-learn and use. More complex functions should be completely optional and possible to only gradually explore. The tool starts with a simple launch screen that asks for the last.fm username. Once a name has been entered necessary downloads are performed automatically (first the listening history, then metadata about the contained songs such as popular tags). Throughout the download, interaction with the visualization is possible and information such as the genre of a song (which is based on user-generated keywords) is gradually filled in once available. At this point, temporal patterns can be easily seen without interacting with the tool. However, the timeline can only show so much information, which is why we again allowed panning and zooming using the mouse. Panning hap-

pens by dragging the background of the timeline and zooming uses the mouse wheel or a pinching gesture on the touchpad. The zoom works only on the x-axis - the scale of the y-axis that shows minutes cannot be changed (to prevent losing oneself in some corner of the timeline).
The upper right corner of the window contains a simple text box that allows searching for songs and filtering for criteria (see figure 5.9). Goal B was to gain more intimate insight into the patterns of the *item* dimension, but to do that without interfering with the visualization we added filtering capabilities. This filterbox supports both beginners and computer-savvy experts (goal E), as it allows various queries without a bloated interface. As an unobtrusive text box it resembles other text boxes on the computer and invites to directly enter queries (e.g., an artist or song name). In case of ambiguities (e.g., a term that could mean an artist or a song) a little arrow next to it allows selecting the intended meaning. By pressing the Enter key, all songs that do not conform to the description are greyed out, leaving only a subset of all songs and its relevant patterns (goal B). While this functionality is easy to discover, the filterbox also allows more complex queries: By entering several terms, all songs conforming to any of them are highlighted. The filterbox also "understands dates ('9h', 'July') and periods of time ('Mo- Fr' for weekdays, 'Jun-Aug' for summer months). Terms can also be negated, to display all items that do not contain them" [ibid]. This functionality allows exploring hypotheses beyond the predefined time scale.
LastHistory's focus on time as the main dimension of the visualization suffers from the

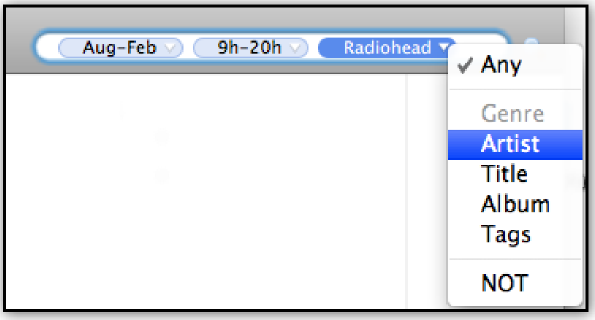

Figure 5.9 The filterbox allows queries of varying complexity. Simple song and artist searches but also concatenated temporal queries are possible. Ambiguities are resolved by selecting the intended meaning (Source: [13]).

same problem that *Strings* has: Other instances of a song cannot be seen at a glance and are spread throughout the timeline. While the always-on arc-based highlighting in *Strings* allowed seeing at least some of these connections, it was more of an aesthetic than an analytical device. We wanted to include this information in *LastHistory* (goal B), but decided to bind it to interaction. By hovering above a song, all other instances of it throughout the history are connected with red curves (see figure 5.10). The curved lines (instead of straight ones) prevent overlap and are easier to follow with the eye [174]. We also used red as a color that had not been used to encode a genre. The arc highlighting is

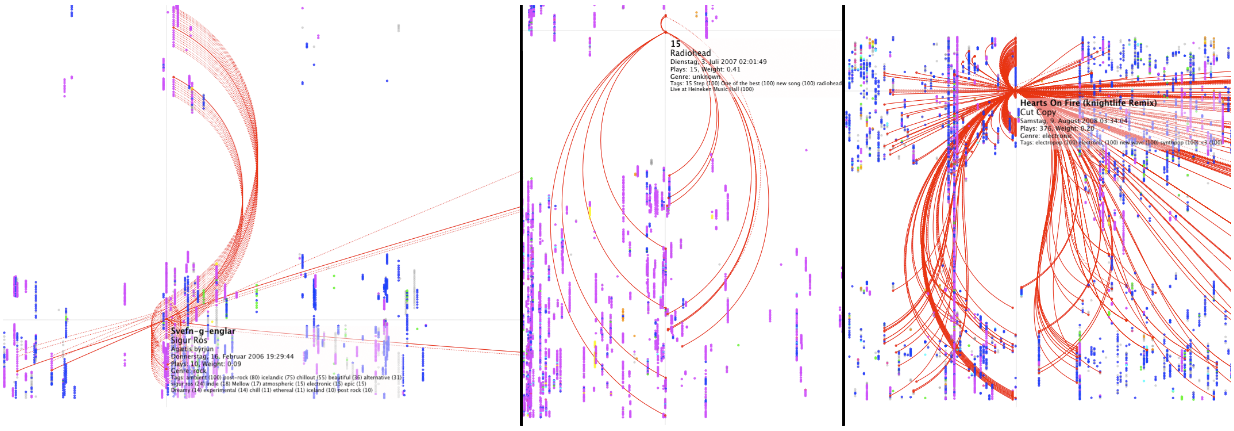

Figure 5.10 Other instances of a song can be highlighted by hovering above one of them. Preceding and succeeding songs are also shown which allows seeing longer repeating sequences. Examples for album listening (left), multiple instances on a single day (middle), and intensive repetitions of the same song (right) (Source: [13]).

not restricted to single songs, however. The tool also examines the immediate neighborhood of the hovered song and draws (dashed) arcs to songs that appear before and after the main song in other sequences. Thus, repeating albums and playlists can be seen and subsequently, if a song is usually listened to in isolation or with common surroundings. Also, the algorithm accepts up to four skipped songs between the sequences and thus allows for some variation in listening.
Another thing that happens when hovering above a song is the display of detail information about it. A transparent text box shows metadata such as artist, title and album, the exact date, the number of total plays, the genre and all user-generated tags from last.fm and their weight.
The *personal* mode (goals C and E), that is supposed to allow for reminiscing and support the owner of the listening history in working with it, provides more information and functionality than the *analysis* mode. Switching between the two modes happens by pressing a toggle button in the upper left corner. One change is that by default *song weighting* is taken into account. In *LastHistory*, the song weight is the importance of a song within a certain time period. This importance is defined by the following formula:

$$w_h = \frac{|t(h)^P|}{|t(h)|} \cdot \frac{|t(h)|}{|t|_{max}} \cdot m \quad \text{with } m = (1 - \frac{M}{2} + \frac{d(h)}{d} \cdot M)$$

"where $|t(h)|$ are the total history entries for track t, $|t(h)^P|$ the number of history entries in the time period P surrounding h, $|t|_{max}$ the maximum number of history entries for any track, d the time interval between first and last entry in the listening history, $d(h)$ the time interval between first history entry and entry h and M a constant describing the influence of m on the weight. $|t(h)|$ does not influence the result, but we chose to show the extended version of the formula to make its derivation easier to explain" [13]. The first fraction contains the importance of a song in a P time window, while the second one

contains its playcount relative to the most popular song (t_{max}). The m factor increases the importance of younger songs, while reducing that of older ones (that have an advantage through their higher number of playcounts).
The song weight determines the size of a track's circle and enlarges more popular ones, while shrinking the other ones (with a minimum size of 1 pixel). This allows seeing popular tracks at a glance and can also be enabled through the menu in the *analysis* mode.
Another change in the *personal* mode is the inclusion of memory triggers. As discussed above (see section 3.5.2), memory triggers are more helpful when including personally created items due to the generation effect. Therefore, *LastHistory* automatically acquires photos and calendar entries from the computer's harddisk and adds them to the timeline, in a separate section beneath the songs (see figure 5.11). Photos are taken from Apple's *iPhoto* application that organizes all photos into events and provides a representative for each event. These events are shown at the corresponding horizontal position. Their size is determined by the number of songs in the listening history at that period, resulting in larger photos for sections of more intense listening. By hovering above the representative other photos from the event are shown. The cursor's position relative to the left border of the photo determines which photo from the event is made visible, with the first corresponding to the left border, the last one to the right and the other photos to positions in between (this interaction is similar to the one in iPhoto). For calendar

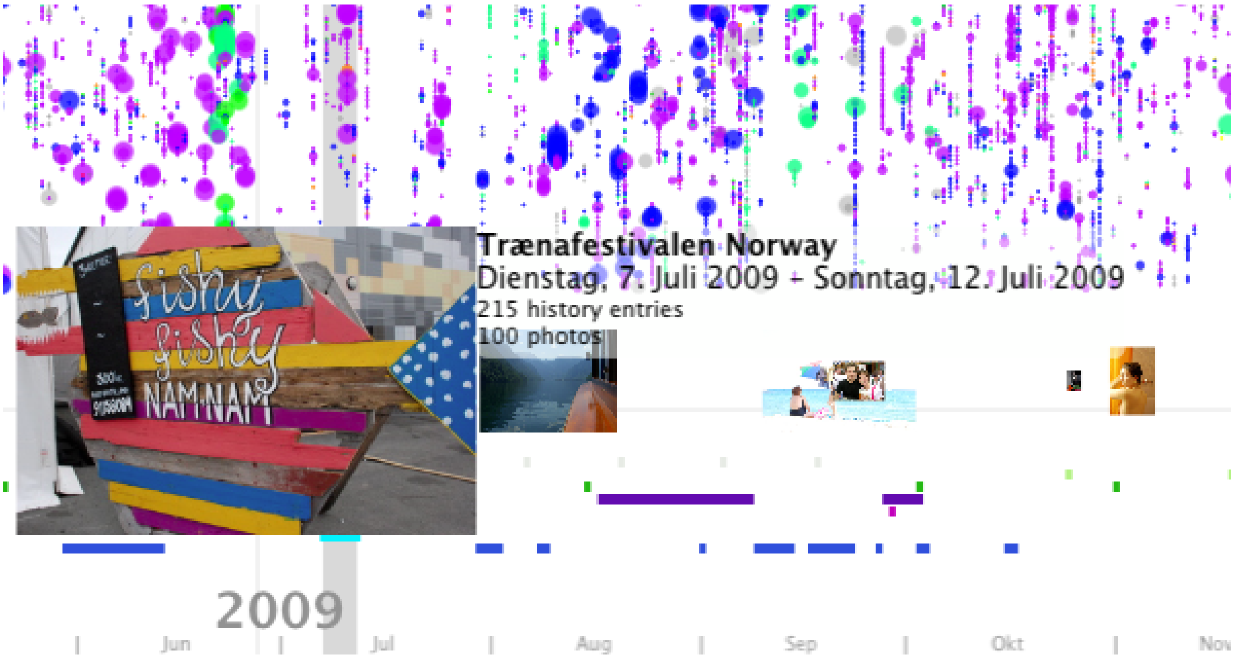

Figure 5.11 In the personal mode, memory triggers in the form of photos and calendar entries are included into the timeline. Additionally, the relative importance of a song in a time period determines its size (Source: [13]).

entries, only all-day events are shown, as we figured that the jumble of everyday duties would only marginally contribute to the understanding or work as a memory trigger.

All-day events contain longer happenings such as business trips of vacations. Hovering above a photo or a calendar event shows additional information in a text box.
Finally, in support for goal D, *LastHistory* also allows listening to music. By clicking on a song node it is played from the local music library. In *analysis* mode, the playback continues with the next song in the listening history. In *personal* mode, songs are arranged by their importance, playing songs with the highest weights first. Minimal controls at the top of the window allow skipping songs or stopping playback. Drawing a rectangle around a set of songs creates a playlist out of them. And by clicking on a photo, a full-screen slide show of all the photos from the event is started and songs from the respective time are played.

5.2.3 Evaluation

The best way to evaluate a visualization tool is still on debate (cf. [131]) and a quantitative analysis for measuring insight is still lacking. We decided to evaluate *LastHistory* in two ways: Our main evaluation would be a short questionnaire for people who downloaded the application from the internet and used it. This way, our participants would represent our intended real-world population and as the survey should only appear after use, they were not aware of being in a study situation when playing with the tool. This approach, however, would only give us superficial information, as the questionnaire necessarily had to be short and would not give us the chance to ask follow-up questions.
Therefore, we decided to do a small lab study exploring the *personal* mode of *LastHistory*. Our participants should use the tool to analyze their listening histories and think aloud.

Case Studies

For our study (cf. [13]) we recruited four participants from our university (21 to 27 years, avg. 24.25, three male) who had been using Last.fm for at least 3.5 years. Their listening histories went from 23,000 to 78,000 (avg. 45,000) items. We introduced them to the tool and gave them tasks to explore available functionality (e.g., "'Use the search feature to find your favorite genres and artists'" [13]). They could then explore their histories on their own. We also videotaped the sessions and asked them to think-aloud.
The results from these case studies were interesting as they filled the abstract listening histories with life (cf. [13]). One participant had consumed a lot of music on weekdays when in school, which ended as soon as she took up a job (the pattern came back for one week on sick leave). One participant had a default playlist to listen to in the morning which appeared every day at the same time. Participants found several gaps within the data, caused by vacation or forgetting to install the Audioscrobbler software. One participant had an increase in scrobbles after taking up tracking music on his mobile player. Songs with a higher weight were popular targets for participants and caused them to rediscover old favorites or remember concerts they went to (and listened to

more music by the respective artist before and after). In one extreme case, the same song was repeated over a hundred times, which the owner of listening history explained with the break-up with his girlfriend. The visualizations were also used to adjust one's self-perception: One participant had to debunk his being a fan of the Red Hot Chili Peppers after discovering that he had barely listened to them in the last two years. The color-coding was appreciated, as it enabled discovering changes in taste (e.g., from hip hop to electronica and *"hardcore-phases"* with predominantly metal music), variety in taste ("*'2009 is color-wise more mixed than 2008'*" [13]), and boasting with listening to music without user-generated tags (grey circles became a badge for exotic taste).
All in all, the discussions with the participants let us see the stories and people behind the histories and we found that it was easy for them to explain some decisions simply by remembering their life at that point in time. The memory triggers were therefore much appreciated. We also found that the tool was easy to use and understand, even without much introduction and we hoped that it would be just as easy for a wider population.

Online Questionnaire

Our main evaluation of *LastHistory* was based on an online questionnaire (cf. [13]). We made the tool available for download and provided a short video introduction on how to use it. We also added a pop-up that appeared when quitting the application after at least 15 minutes of use that asked to fill out the questionnaire. The pop-up only appeared once without nagging everytime *LastHistory* was shut down. The questionnaire was kept short to entice people to fill it out completely and not give up on the second page. It contained questions about demographics, if photos and calendar entries were available, how the application was used, which of the two modes they liked and what they had found, and how to improve the application.
After making the application available, several blogs and news sites picked it up (e.g., Newsweek[2], Lifehacker[3], Music Machinery[4]) which helped us in creating a broad user base. In October 2010 when writing the paper ([13]) and half a year after release, we had reached 5,000 downloads and 243 completely filled out questionnaires. This low feedback rate of around 5% came not entirely unexpected as we estimated that around half of the downloads would either result in only trying the tool shortly or not at all. The length of the questionnaire seemed to be appropriate, as we had more completed (243) than non-completed (155) questionnaires.
A first set of results concerned demographic information. Asking for gender, age and occupation allowed us to learn more about our population. Based on the general last.fm population, we expected a bias towards younger and male people but were surprised by the actual result of 95.1% male to 4.9% female. They were ranging in age from 16 to 67 years (avg. 27.2) and had mostly a background in technology or academia (56%). As

[2] http://www.thedailybeast.com/newsweek/blogs/techtonic-shifts/2010/03/04/lasthistory-mashes-up-your-music-and-photo-timelines.html

[3] http://lifehacker.com/5493059/lasthistory-graphically-visualizes-your-lastfm-history-through-time

[4] http://musicmachinery.com/2010/02/16/lasthistory-visualizing-last-fm-listening-histories/

for the general use, nearly everybody (99.2%) used *LastHistory* for visualizing a listening history, with 97.1% their own and 4.5% an external one. In 37% of all cases photos, and in 18.5% calendar entries were available. 46.9% said that they used the tool for reminiscing and 8.6% for storytelling. The *analysis* mode was more popular (75.7% found it useful) than the *personal* mode (51.8%). The reason for that, we discovered when analyzing the data, was the lack of available memory triggers: People who had both photos and calendars available found it similarly useful (75.8%) while people without this data disliked it. People were also in general happy with how easy to learn (75.4%) and to use (75.7%) *LastHistory* was.

We were especially interested in the free form fields that the questionnaire contained and where participants were asked what interesting aspects they learned about their listening histories when looking at them in *LastHistory*. Similar to the ones in the case studies, the answers conformed with the listening factors (see section 2.4). People often understood what had happened at a certain point in time simply by looking at the visualization: "*'During the period where I wasn't often listening to much music during the night time, when I looked at the few instances of music playing in the early morning, I found that in many cases I could recall quite clearly the night that that happened, from the time and the tracks that were playing'*" [147]. "*'My listening time periods have changed, and I can identify those time periods with what's going on in my life'*" [ibid]. One participant lauded the *personal* mode for "*'Being able to connect people and places with my listening habits'*" [13]. Goal A, finding temporal patterns, was supported well, as many of the answers mentioned them: "*'I rarely listen to music between the hours of 9-11 a.m., even on weekends', 'Patterns across different academic terms', '[I] noted the ... commuting listening pattern', 'Those ruts where you get stuck in listening to one particular song'*" [ibid]. The secondary goal of item patterns also showed, as some participants explicitly talked about musical taste or the combination of time and items: "*'How often I listen to Grizzly Bear!', 'I listen to Aerosmith around 7pm quite a lot of the time'*" [ibid], "*'I go through waves of listening to genres', 'I stick to the same genre for a few weeks then move on'*" [147]. Participants also discovered problems with the data ("*'That I listened to music for 4 straight days. Apparently my computer was on and played the music without my knowledge', 'Huge gaps in my listening when I wasn't scrobbling'*" [ibid]). As mentioned above, photos and calendar events were appreciated if available ("*'Clicking on a photo gallery and listening to what I was listening to at the time was very powerful'*" [ibid]).

In general, the online evaluation of *LastHistory* first of all showed us that there was actually a need for such lifelogging analysis tools. Second, we were able to learn about the characteristics of people who used it and also, how many of them had data for memory triggers available (even though some people complained that it was not their lack of photos or calendar entries but just the support for the wrong tool). Third, we found that the listening factors existed and that people were able to understand the listening context both from the visualization and the memory triggers and relive the depicted time of their lives using *LastHistory*.

5.2.4 Discussion

LastHistory is a successful example for more complex listening history visualizations. Its evaluation showed that it supported a need and helped people learn about themselves. The visualization concepts that we used could also be applied to other media (cf. [13] and section 7.2), but should take into account the differing aspects of consumption: Books and movies, for example, take in general more time to be read or watched and are also not repeated as often as a piece of music. Relying on a suitable abstraction towards genres or "artists" (maybe again by color coding) could help uncover patterns. The timeline in *LastHistory* also proved flexible enough that other media might just be incorporated, to provide a more overarching picture of one's media consumption.

Regarding the implementation, compared to *Strings*, *LastHistory*'s was more low-level and relied on the Mac OSX's Core Graphics framework. We also relied on the availability of iTunes (song files), iPhoto (photos), and iCal (calendar events) for music and contextual information, which meant that *LastHistory* was Mac-only. Still, this made it convenient to use (it automatically loaded information from the harddisk without having to set photo paths etc.) and much faster than *Strings*: Histories with less than 100,000 songs caused no performance problems on a regular Macbook Pro (2009 edition). The main bottleneck was the download of the history and accompanying metadata from last.fm, which took around 45 minutes for a 25,000 item history. Some feedback from the questionnaire also mentioned that ("*'It is very sluggish for me'*, *'Speed it up!'*" [13]).

Even though photo and calendar events had proven to work in uncovering context from memories, other contextual data sources would probably be helpful and easy to integrate with the timeline (see section 3.5.2). Status updates or blog posts and other personally created items should help.

Another thing we observed after *LastHistory* had been released was that people started publishing screenshots of it on their blogs or on flickr (e.g., here[5]) and using them to explain their listening behavior (similar to the various last.fm badges, see section 4.1.2). We had not expected that, but it made sense in retrospect especially in an environment as focused on identity management as last.fm. Future lifelogging visualization tools should contain some form of sharing functionality and maybe even support annotations, so people can explain their behavior without having to resort to manipulating the data (cf. [153]).

In any case, this sharing of the static visualizations showed that our approach of having an immediate benefit without requiring interaction was the right way to go. Also, the gradual discovering of optional functions gave people the feeling that the tool was easy to use (75.7% of the online participants thought so) and did not overwhelm them.

[5] http://www.flickr.com/photos/phnk/4448476537/

5.3 LoomFM

So far, the presented visualization prototypes were only concerned with single listening histories. While *Strings* and *LastHistory* explored the *time* dimension in depth, the *listeners* dimension was completely left out. I first wanted to get a sense for the difficulties of uncovering patterns in a single listening history before approaching the problems of additionally supporting patterns between multiple histories. Using the experiences from the first two prototypes I developed *LoomFM* (see figure 5.12 and [33]). The main idea behind the application is showing similarities and differences between two listening histories. One of them will usually be owned by the analyst, while the other one will be that of a friend or family member (a stranger's listening history is usually uninteresting). For this visualization, I again put an emphasis on *time*, but the comparison aspect played the central role.

5.3.1 Population and goals

Similar to the above visualizations, the usual last.fm profile owner was the intended main user population, which again resulted in only fuzzy descriptions of background knowledge on visualizations or music. The tool should be easy to use and provide immediate benefits. The goals that *LoomFM* should enable were:

A: Compare overarching patterns between the two histories. The main goal of *LoomFM* is comparing histories, and this should foremost happen on an abstract level. Finding general trends and quirks is more important than having a detailed insight into the song-to-song behavior. Disputes such as who heard an artist first or who is listening to more or less music of a certain genre should be resolvable using the tool.

B: Finding temporal patterns. This goal is similar to the ones in the last two visualizations, as having a clear localization of song events in time helps in understanding listening decisions. Also, having time-based information allows assigning patterns to certain time frames.

C: Compare histories on a more detailed level. Even though an overarching view of the two compared histories might be helpful, sometimes having a deeper understanding is necessary. Therefore, *LoomFM* should support the full spectrum of details regarding music and also allow comparisons on lower levels, down to the song level.

In addition to the goals, I planned to keep the interaction minimal in *LoomFM* and provide as much information using the static visualization as possible.

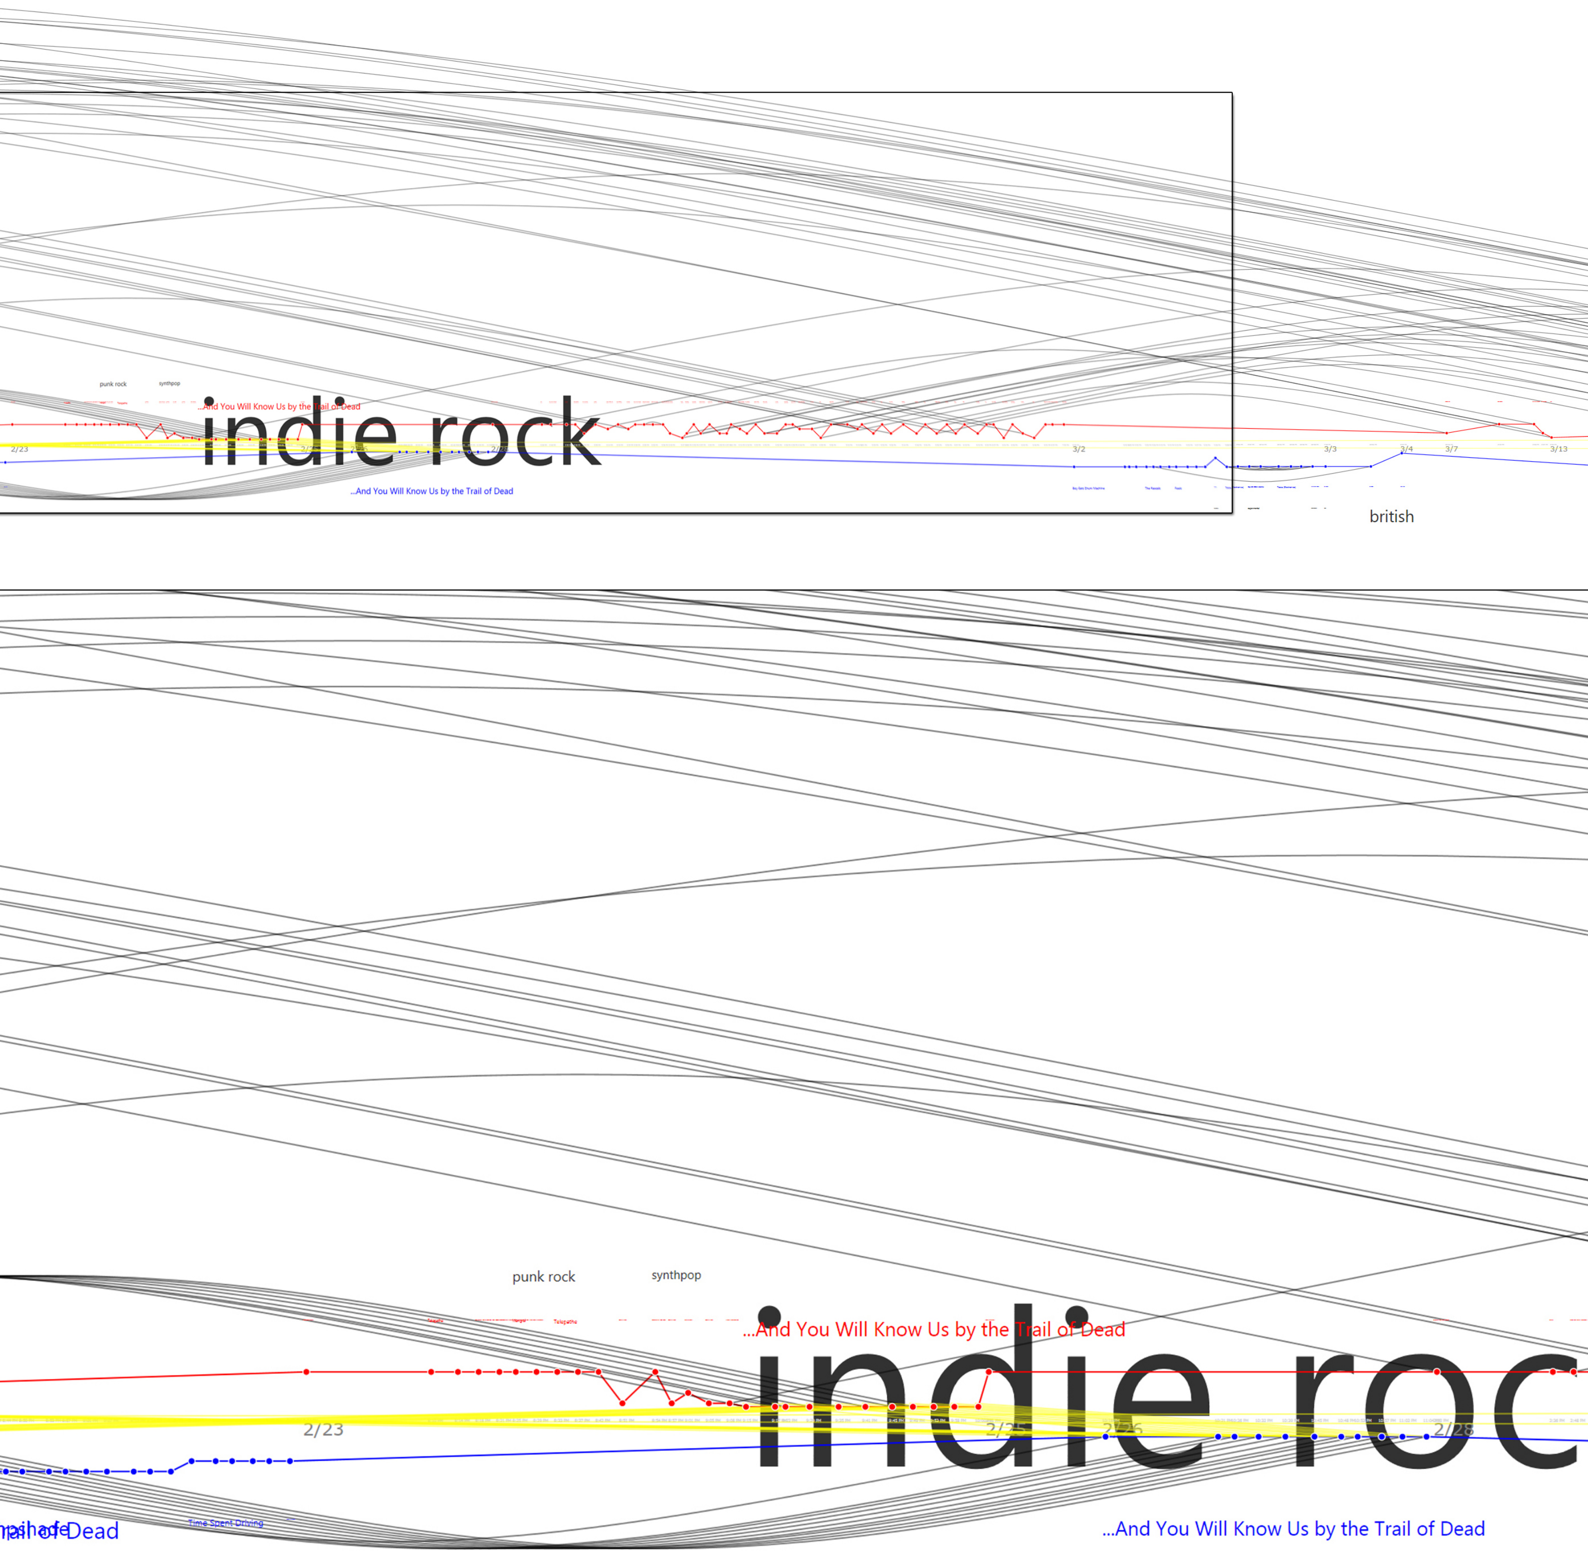
punk rock
synthpop
...And You Will Know Us by the Trail of Dead
indie rock
...And You Will Know Us by the Trail of Dead
british
punk rock
synthpop
...And You Will Know Us by the Trail of Dead
indie roc
2/23
Time Spent Driving
...And You Will Know Us by the Trail of Dead

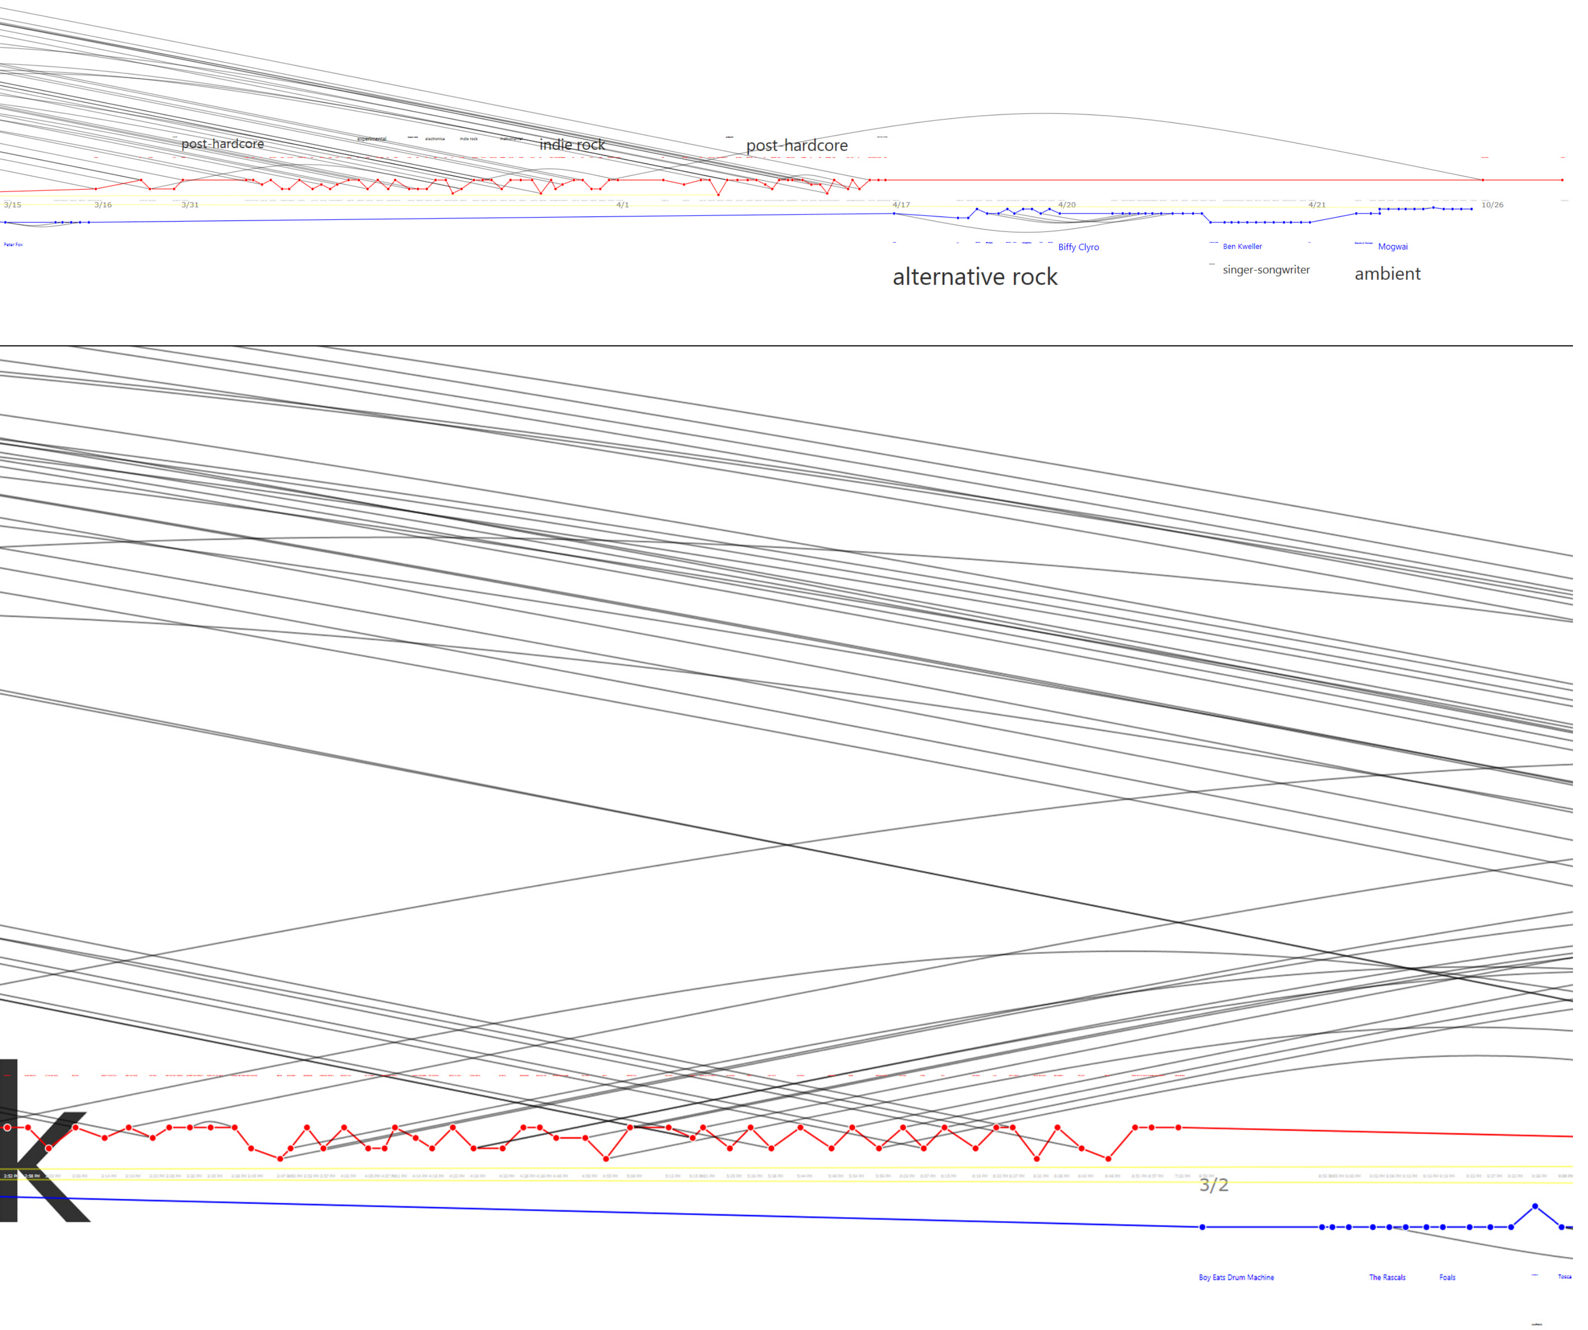

Figure 5.12 *LoomFM* compares two listening histories on the song level. A horizontal timeline shows temporal patterns (top). The enlarged section (bottom) shows single songs and overarching artists and genres.

5.3.2 Design

The main design element in *LoomFM* is the timeline. Even though goal A suggested having overarching patterns as the central aspect of the visualization, showing temporal patterns and aligning the song events with one's own life were easiest to realize by sorting them along a timeline. In *LoomFM*'s case, the timeline is horizontal and only one-dimensional - all songs are sorted by their dates. The vertical dimension is used to display the similarity between songs: While the timeline is in the center of the screen, the first listening history is above and the second one below it. Each song is represented by a small circle and the two histories are color-coded (red and blue). The relatedness between songs (goal C) is expressed by their closeness to the central timeline: The "closer" a song from one history is to the other musically, the closer it moves towards the center. Closeness ranges from neither the song nor the artist ever appear in the other history, to the same song is played at the same time in both histories (which is rare). Based on this scale, the two histories look like fever charts of similarity that show the ups and downs in taste (see figure 5.12 (bottom)).
Naturally, listening histories contain more songs than can fit on screen, so some form of navigation of the timeline had to be included. Similar to *Strings*, clicking and dragging the mouse pans the canvas and the mouse wheel zooms in or out. Still, this is almost all the available interaction, as I tried to keep the visualization in general very static. A second form of interaction available is hovering over song items or arcs (see below), which provides more information, an artist image, and at what time it appears in which history (see figure 5.13).
This minimal interaction still effectively support the main goal A of uncovering overarching patterns between the histories. In general, the level not only of graphical but also of informational detail is bound to the zoom level: The closer one zooms in, the more detailed information becomes available (similar to a semantic zoom [130], but only bound to the graphical representation). This informational content comes in the form of labels for artists and genres that follow the rule that the more important an artist or a genre is at a certain point in time, the larger its label becomes. The size of these labels therefore depends on the number of consecutive songs with the same artist or genre (see, for example, 'indie' in figure 5.12 (bottom)). Based on this, labels of often-heard genres and artists grow in size and can be seen even on low zoom levels, while one-hit-wonders only appear after zooming in closely.
Another form of overarching patterns are repeating songs (goal A & C). All timeline-based visualizations are plagued by the spreading of multiple instances of the same songs over the canvas. Similar to *Strings* and arc diagrams [175], I decided to use arcs in *LoomFM* that connect identical songs. This not only happens within one history but also between histories. The size of the arc grows with the distance between the two song instances, which means that the height of the arcs over a listening history depicts its owner's tendency for repetitions. Small or no arcs mean only infrequent repetitions of recently heard songs, while high arcs show repetitions even of older songs. The red history in figure 5.12, for example, shows high arcs and frequent repetitions, while the

blue history contains almost no arcs. As these arcs also connect songs in different histories they can show overlaps in taste. And they also highlight sequences: Figure 5.12 (bottom) shows the release of a new *...And You Will Know Us By The Trail Of Dead* album (very left) and the subsequent in-order listening by both red and purple.

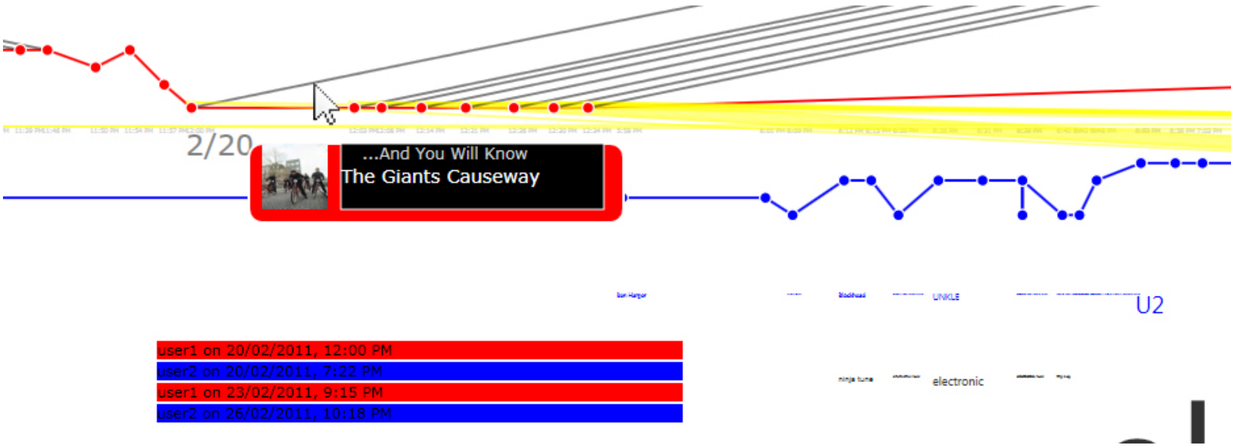

Figure 5.13 Hovering over a song node or arc shows detailed information about the song and the distribution of instances throughout the histories.

5.3.3 Discussion

LoomFM is similar to *Strings* in its simplicity and playfulness. It supports one main task (namely the comparison of the two histories) and does so with minimal interaction. The learning curve is non-existent and the tool's one-off task requires no intense training. The connections between the listening histories can be explored simply by zooming and panning. I decided to skip an extensive evaluation for the visualization due to its restriction to this simplicity.

The implementation in C# with Microsoft's WPF framework ensures a certain speed and responsiveness, even with hundreds of songs. Still, aside from the implementation, scalability is a problem as, for example, arcs between the histories only have the small space around the timeline to be visible, which generally results in one long stretch of overlapping arcs.

Also the visualization is less flexible than, for example, *LastHistory*, but can be quickly picked up and used. Adding a secondary history has the potential to increase the complexity not only of the tasks and patterns available but also of the interface (see next section for an example).

5.4 LastLoop

What all the above visualizations have in common is that they focus on one specific dimension of the design space: Both *Strings* and *LastHistory* supported mainly *time* as a data dimension, while *LoomFM* was also time-centric but put its emphasis on the *listener* dimension. Focussing on one explicit dimension makes it easier to adjust the necessary visualization and interaction to uncover patterns within it. However, overarching patterns between all three dimensions are practically impossible to find (even though memories can help, cf. *LastHistory*). For the last project that used listening histories for analysis and reminiscing purposes, we[6] wanted to support all three dimensions in an equal measure. On a very abstract level, our goal was to create a crossover between *LoomFM* and *LastHistory*, i.e., a tool that would by easy to use and support multiple listeners, but also allow complex insights and provide corresponding functionality. We called the resulting tool *LastLoop* (see figure 5.14). *LastLoop* should not only provide more functionality than the other tools, but also be more flexible in its deployment (we did not want to have a Windows- or Mac-only tool). It should also be evaluated with a sample of real-world last.fm profile owners (similar to *LastHistory*).

Figure 5.14 *LastLoop* shows all three data dimension of *time*, *items*, and *listeners* and allows finding patterns within the data.

[6] *LastLoop* was part of Roman Graebsch's Bachelor thesis [58].

5.4.1 Population and goals

As our intention was to support complex analyses with *LastLoop*, the user population was very broad: It contained not only the varied last.fm population, but also possibly analysts with research backgrounds. The lack of proper tools for analyzing listening histories (cf. section 5.2.1) also encompasses tools for comparing multiple histories. Also, as the main goal of the visualization was to provide a suitable analysis tool for multiple histories, it was clear that not all of them would be owned and known in-depth by the person in front of the computer. Therefore, we did not include memory triggers as in *LastHistory*, as they would have made the interface more complex and been only of minimal use. In any case, we could again not be sure about musical and visualization background and tried to keep the interface as accessible as possible for people with different backgrounds.
With *LastLoop* planned as a mixture between a complex single-listening history visualization tool and a more playful tool for multiple histories, the general analysis goals were also a mixture of theirs:

A: Find temporal patterns within multiple histories. This goal was similar to the one for *LastHistory* except for its extension towards multiple histories. *LoomFM* had finding overarching song patterns as its main goal and thus its approach was more informal. Just as with additional data sources, *time* is the one common element for multiple listening histories that might otherwise be very different in their musical content. Therefore, having *time* as the one overarching dimension would make the most sense.

B: Find musical similarities and differences. While the first goal made sure that some form of comparison would always be available, in light of the social nature of music (see section 2.1.3) a comparison along the *items* level could bring far more interesting insights. Such comparisons could happen along the whole scale of music from single songs to genres and should be reflected accordingly in the interface. These different levels of detail would enable comparisons of different granularity.

C: Focus on certain aspects or sections of the histories. In the previous visualizations the data set was generally restricted with only one or, at most, two histories being available. Especially the experiences with *LastHistory* had shown us that some type of filtering was helpful and given that *LastLoop* could work with not only two but much more histories, focussing the analysis on some aspect would be very important. This filter should work on either of the three data dimensions or a combination of them.

D: Not be frustrated by the application. Given the complexity of the subject matter, we expected a high level of potential for frustration. In addition to the already difficult use case, we did not want to overburden analysts with problems of the tool. Therefore, as much as possible should happen automatically and the interface should be, as always, easy to use and provide immediate benefit.

The resulting visualization should be similar to the *analysis* mode in *LastHistory* with its lack of additional data sources and should represent a superset of its functionality: In the best case, loading only a single history in *LastLoop* should have allowed the same insights as *LastHistory*'s *analysis* mode. In addition to that, however, it should also support a larger number of histories without sacrificing ease-of-use or possible insights.

5.4.2 Design

As we were well aware that we could not produce a high level of fidelity in insight for an arbitrary number of listening histories (e.g., the whole last.fm population), the design of *LastLoop* always aimed at a number between two and five of them. This would mean that the tool could provide insight for small groups of people, friends and family members, and do it on a detailed level. Yet, the tool itself was not restricted to five histories, so in practice an arbitrary number of them could be added.
The potential difference between various histories meant that *time* was their main dimension for comparison (goal A). We therefore again used a timeline as the central visualization metaphor. To allow for more detailed insight into the temporal patterns, we also adopted the two-dimensional timelines from *LastHistory* that showed date on the horizontal and time of day on the vertical axis. Simply taking one timeline and adding all song instances from all histories to it would have led to a confusing display of multiple overlapping songs from different sources. Even when using color coding to distinguish between the different histories, the abundance of music throughout a common listening history would in any case have created a lot of visual clutter and complexity. We therefore took the same approach as with the additional memory trigger data in *LastHistory*: The available vertical screen space is equally distributed among the number of histories and each of them has its own two-dimensional timeline. While this prevents comparisons regarding the time of day when music was consumed (visual comparison between adjacent bands is difficult), it nevertheless allows them for the day-level. Each song was again represented as a small circle, which allowed estimating differences in intensity and musical involvement at a glance. The same patterns that were available in *LastHistory*'s temporal display were also provided by *LastLoop*, thus creating the superset in functionality that was mentioned above. Through the visual adjacency, these patterns could however be immediately compared. Different histories can potentially be from different countries and the shift in timezones (see section 5.2.2) between them is also easy to see.
As all user attribution is encoded through spatial means, color is available for providing musical information. We encode the genre of a song using static colors, but with different attributions than the ones from *LastHistory*. We extended the color scale a little and chose 9 different genres (rock: dark green, metal: light green, jazz: dark blue, funk: purple, classical: light blue, electro: orange, pop: brown, hip hop: pink, blues: rosy) and white for an unknown genre. Their functionality is equal to *LastHistory* and provides insight of musical similarity, thus supporting goal B.

The common problem of song distribution throughout the timeline intensifies with the introduction of multiple histories. In a scenario with up to five histories the same song might appear not only within one of them but potentially all of them and multiple times. Therefore, we had to find a way to show these connections without overburdening the visualization or creating a completely indecipherable display.
Our approach relies on arcs that connect songs but also heavily on ways to filter and adjust the display of these arcs. For one, arcs instead of lines prevent clutter and overlap, as discussed in section 5.2.2. But it is also important to distinguish between connections within one history or between histories. Therefore, intra-history connections are drawn in red, while inter-history connections are yellow arcs. Additionally, arcs for internal connections appear within the visualization plane, while arcs between histories have a slight 3D effect as if they were spanning above the visualization (see figure 5.15a).
 The interaction also plays an important role in keeping the amount of information

Song instances are connected with arcs. (b) Each listening history has a control panel.

Figure 5.15 Intra- and inter-history arcs and control panels in *LastLoop*.

manageable. Navigating the timeline happens with a slider widget at the bottom of the application, that controls the position and zoom level of the timelines (all visible timelines are coordinated and are bound to the same dates). Dragging the timeslider navigates within time and grabbing its left or right ends allows adjusting its width (and thus the zoom level). Zooming is also possible with the mouse wheel.
Each listening history has next to it a control panel (see figure 5.15b) that allows adjusting the appearance of arcs. By default, arcs only appear when hovering above a song node (they can be fixed with a left-click and released with a right-click). However, it is also possible to always display all intra-history arcs (as in *Strings*) - which can be overwhelming - or selecting a certain time section and only show arcs that either originate from there or originate and end within the section. All inter-history arcs cannot be shown at the same time, but they are displayed for a specific song when hovering above it or can be switched on for arcs originating from the selected section. Choosing

a specific time section allows reducing the number of arcs while not having to hover from song to song to see all relevant connections (goal C). The control panels also shows how many other instances of the currently selected song are visible and available in total (red number) and how many arcs are extruding from the current selection area (yellow number).
To focus on other aspects than time, controls at the bottom right of the application can be used to filter for genres or searching directly for artists, songs, and titles (see figure 5.16). Sought-after items are highlighted (in turquoise) while filtered genres are hidden. Also, the respective control panels show how many of the sought-after songs are visible and available (turquoise number) and how many songs the listening history contains in total (white number).
Using this in combination with the timeslider, all three dimensions can be adjusted and

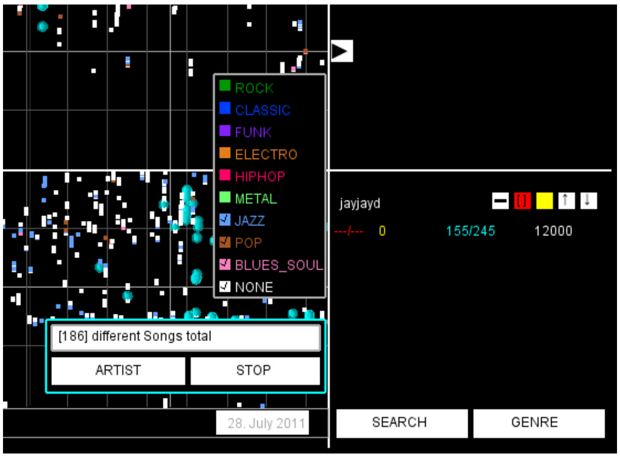

Figure 5.16 The two controls at the bottom right of the application allow filtering genres and searching directly for artists, albums, and songs.

the analyst can focus on specific questions when looking at the histories.
One piece of feedback that we often got for *LastHistory* was the frustration that the tool had been Mac-only. We accepted this trade-off to make it as conveniently to use as possible (we could seamlessly integrate contextual data without requiring any program dialogs or settings), but were in a different situation for *LastLoop* as there was no equivalent to the *personal* mode. Therefore, and to support goal D, we wanted to make it purely web-based. We chose the processing[7] framework for that, a Java-derivative aimed at visual designers that provides a very developer-friendly way to create graphical output. Additionally, processing also allows creating Java applets, that can be launched from the browser without any installation and even support offline cache (this was helpful as we did not want to strain last.fm's servers by repeatedly downloading the same histories). After starting *LastLoop*, profile names can be entered in the text box in the upper right corner of the application and corresponding data is loaded automatically (interaction is possible the whole time while the data gradually fills the visualization).

[7] http://www.processing.org

5.4.3 Evaluation

For evaluating *LastLoop* we had a similar approach as in *LastHistory*: We made the application available online and added a questionnaire. Due to the different implementation, the experience was a bit different: When visiting the lastloop website[8] and clicking on 'LastLoop App', the visualization's Java applet was immediately launched in a new window. The old window had links to a video explaining the application and to the questionnaire. It contained questions about involvement with music, computer experience and their enjoyment of the application. In addition to the questionnaire, we also added a button labelled 'Feedback' in the upper left corner of the application and asked people to "report anything interesting [they] found". What the button did was take a screenshot of the current state of the application and send it and the comment that people left to us.
Unfortunately, the amount of feedback we received was only small, compared to *LastHistory*. After two months we had 20 complete responses. Little surprising, almost all participants listened to music daily (95%) and were experienced with computers (80% answered 4 or above on a five-point Likert scale). We used 18 five-point Likert scales for asking about general aspects of *LastLoop* (e.g., "*'I learned something about the way I listen to music'*, or *'Overall, I am satisfied with LastLoop'*" [58]). People were generally neutral or slightly positive about the different aspects and all questions had a median answer of 3.
The free text questions about insights found and ways to improve the application had more interesting results. Insights fell into two categories: For one, people learned about their own music listening habits (only two participants had used other visualizations before) and mentioned, for example, that "*The way I listened to music was deeply related to the things that were happening in my relationship*" or "*I listen to music mostly for about 4-5 hours without a break*". On the other hand, people also used *LastLoop* as intended and found patterns between histories: "*Inter user connections showed interesting facts: when did the other user hear my favourite song, have there been many connections lately*", "*That one user is also listening to a very infamous band, from the 70th [sic]*" (all participants comments from [58]). The relatively small number of responses, however, did not let us derive general trends from these answers and besides expected responses we did not find anything new.

5.4.4 Discussion

It was interesting to compare the interest for *LastHistory* with the lack of it for *LastLoop*. While both applications had similar goals and practically identical user populations, the former gave us a solid number of responses, while the latter brought only very little feedback. We saw several reasons for that:

[8] http://www.lastloop.de

First, *LastHistory* had the big advantage of being picked up by mainstream media very quickly and thus generating a large number of downloads. Even though it was Mac-only, thousands of people used it which increased the chances of receiving feedback. *LastLoop*'s integration into the browser should have been more convenient, but even there people encountered problems when they either had no Java runtime environment installed or a wrong version of it. When developing LastLoop in Summer of 2010, native browser-technologies had not been as advanced as they are today so Java was the best solution, but nowadays choosing HTML5 and Javascript instead would be better.
While providing a short tutorial video was enough to explain *LastHistory*, *LastLoop* proved to be too complex for this approach. Adding more videos, documents explaining the functions or giving tutorials would have helped.
Regarding the tool itself, its main problem was supposedly the complexity of the underlying data and the resulting interface. Adding the *listener* dimension to the visualization multiplied the possible patterns and could be overwhelming. Based on the feedback we also identified other problems that plagued the application. For one, only little help was available within the application and it was difficult to remember the functions of the various buttons. Second, hovering over songs was also only bound to the horizontal position of the mouse cursor instead of both horizontal and vertical, which could be frustrating. This made it difficult to select songs in low zoom levels.
Mainly, *LastLoop* did not provide as much guidance as the other tools and required more exploration. Immediate benefits were available in the form of multiple timelines, but only after discovering how to load listening histories (with a button in the upper right corner). Also, the interface itself with its widgets and buttons looked different from existing ones so people could not rely on their experience (processing does not provide regular UI widgets and they have to be built from scratch) and was immediately available in all its functionality. The interface was therefore a common criticism (participants asked for "*more convenient controls*" or "*a [sic] easier to use interface*"[58]). While *LastLoop* was powerful and allowed complex analyses, it might just have been too much of an expert's tool to be useful for casual usage.

5.5 Summary

Looking at one's music listening history through the lens of a visualization is still a vanity affair. Having little dots represent songs and spanning a timeline does not help in losing weight or becoming a better person. And analyzing such histories is still more about stroking one's ego than making the world a better place.
However, looking at the excitement in people's faces when rediscovering long forgotten favorites or reading about their memories of loved ones and bright days at the beach uncovered only through little colored circles showed us that listening histories were more than just long sequences of songs. They represented parts of personal histories and worked as pointers towards long lost memories. A single song node could be like

a madeleine and potentially hold a whole month of stories. Reminiscing was thus the most powerful use case for listening history visualizations and memory was the most powerful data source. Even the most obscure explanations for listening decisions appeared perfectly logical once the histories' owners had the chance to speak their minds and these explanations were also immediately available. Using personally created items as memory triggers was useful and possible without much effort. Relying on memory triggers might not help in countering all lack of data in lifelogs but should not be underestimated as a data source.
Equally important in this regard is providing a way to map the items to a personal life story and points in time. Our reliance on timelines was also a conscious decision for supporting the analyst's memory. While simple dates might not always be enough to make sense of a history, they do no harm either and together with other song instances they might help in forming an overall picture.
A last lesson that we learned was that trying to address all three data dimensions of *time*, *items*, and *listeners* as we did with *LastLoop* might be too complex to be handled by an interface and casual last.fm profile owners. There seems to be a gap between single and multiple histories, as most people are able to find interesting patterns in their own stories and also enjoy this more. When confronted with the chance to analyze more than one history and compare them, they have more trouble finding a suitable approach or analysis task. Focusing on a certain aspect of a listening history or at least one of the data dimensions is helpful for casual use.
Regarding the presentation and functionality of the interface we were also able to draw important conclusions: The best form of the interface of such a listening history visualization tool is based on its usage context: Analyzing a listening history or using it for reminiscing does not happen on a daily basis which means that it is very important to keep the experience as little frustrating as possible. Also, entry hurdles should be kept low and the interface has to be usable more or less without instructions (as nobody will go through the pains of re-learning the interface every few weeks). Relying on experience with user interfaces and using familiar widgets consistently can help. Also, doing as much as possible automatically (e.g., loading photos and calendar events from iPhoto and iCal in *LastHistory*) lowers the strain on the analyst.
As the tools we created were mostly new and people had no experience with them and we could not be sure about their experiences and background, they had to be flexible and at least look simple.
One helpful concept was that of providing an *immediate benefit* with the least interaction necessary (only entering the last.fm username). The time-centric visualizations were appealing not only in their aesthetics, but also in their clear meaning (even in the two-dimensional case). They immediately gave something back and enticed the owner of the history to dig deeper and start exploring the application. An immediate benefit could also come from the chance to listen to music by clicking on a song (something that most people would try immediately in *LastHistory*). When requiring more input or more obscure interaction before getting something back (as in *LastLoop*), people are more sceptical which changes the complete experience they have with the application.

Providing something useful as early on in the interaction as possible is central for securing a positive attitude.
Once this idea of the application being useful has been established, more complex interactions become possible. However, the interface still has to *look simple* and accessible, so people are not afraid to try and explore it. Providing all functionality at once (see, for example, toolbars in older office software) might help in making interaction faster but also renders it more complex and intimidating. As these visualization tools are not used on a regular basis, making their interaction processes slower is a good trade-off for keeping them accessible and friendly. Complex tasks should be supported but they may require a few clicks more than in the most efficient user interface.

The next chapter shines more light on the use of listening histories for music-related tasks. The main idea behind it is that of the *items* dimension and the second way to look at a listening history: Not as long sequences of items, but as an intricate network of item relationships.

Chapter 6

Listening histories:
Visualizations for playlisting and rediscovery

No more talk about the old days
It's time for something great.

– Thom Yorke - *Atoms for Peace* –

Long-term music listening histories spanning multiple years capture one specific aspect of life. But they can also be interpreted in a different way: When ignoring the *time* dimension, listening histories become expressions of taste. A song that appears only once and is then not even finished was probably less favorable for the listener than another one that has been played more than a hundred times. Simply counting the number of appearances of songs can already be used to create a picture of a listener's taste in music. This approach is used by regular media player software such as iTunes, that counts and displays for each song the number of plays so far and also sorts the table of songs based on this criterium. Such implicit votes about favor are, of course, not as accurate as an explicit expression of it (e.g., Facebook's *Like*-button), but they are still able to show a trend: These votes with your ears mean that even if the listener was more enduring than enjoying a song, it was still more pleasant to listen to it than to skip it.

Going beyond this simple counting can uncover more interesting aspects of personal music listening. A first step is looking at the predecessors and followers of a song in a listening history (preferably within the same listening session), and analyzing these sets. It might be that their contents are not completely random and that there are tracks that have a high probability of appearing before or after the song in question. Such a co-occurance pattern means that there exists some form of relationship between these songs, for whatever reason. This relationship is also transitive, as there might be whole sequences of songs that repeatedly appear throughout a listening history. One can go also further and take not only the immediate predecessors and followers, but the whole graph of songs before and after and determine the relationship based on the distance in sequence. Songs that have been played on the same day then are still closer than those that only appeared in the same month.

The advantage of such analyses is that they are entirely possible without knowing anything about the complex issues of listening context (see section 2.4) and personal stories. The songs might appear in sequence as they had been first heard on a memorable vacation trip or when meeting one's significant other for the first time. The reason might also be much more trite if the songs follow after one another on the same album and the listener has a tendency to listen to whole albums instead of creating playlists. As explained above (see section 4.2.1) relationships between songs are either inherent to them (e.g., similar content, artist, lyrics) or are event-based and created by the listeners themselves. Whatever the reason, the relationship between the songs can be trivially extracted from a list of songs and timestamps alone. And these relationships represent the song network seen through the listener's very personal lens. Even though the system might not know what caused the connection between tracks, automatically creating playlists from such songs is convenient and trivial. *LastHistory*, for example, allowed listening to old sequences from the listening history and this feature often received positive feedback. The playcount can also be used for provoking rediscovery, by playing songs that repeatedly appeared at some point in time and not between then and now.

By looking at the full spectrum of all listening histories, the accuracy of song relationships can be improved. A single listening history necessarily contains characteristics specific for the listener: A pair of tracks might often appear in sequence just for sentimental reasons or out of sheer luck. Taking a larger sample is enough to arrive at the inherent relationships only and ignore the personal and event-based ones. Merging multiple listening histories can thus help in producing a "statistical" picture of these relationships of taste. This is also the basic idea behind collaborative filtering. Co-occurance means some kind of similarity and for a large enough sample the calculated similarity will be sound.

In this chapter, I present three prototypes that address this relationship between songs (the *items* dimension), either in its inherent or event-based forms. When looking at these relationships, time is often not as important as in the lifelogging cases (see previous chapter). The purpose of these tools is more about listening to music, building playlists or rediscovering old favorites than understanding listening decisions and reminiscing (even though there can be an overlap). Also, when using the aggregated relationships, the results are more important than the sources.

Rush is an interaction technique for mobile devices that allows creating a suitable playlist quickly while adjusting the result for intended purpose or current listening context. It is based on pure inherent similarity between songs and the whole last.fm data.

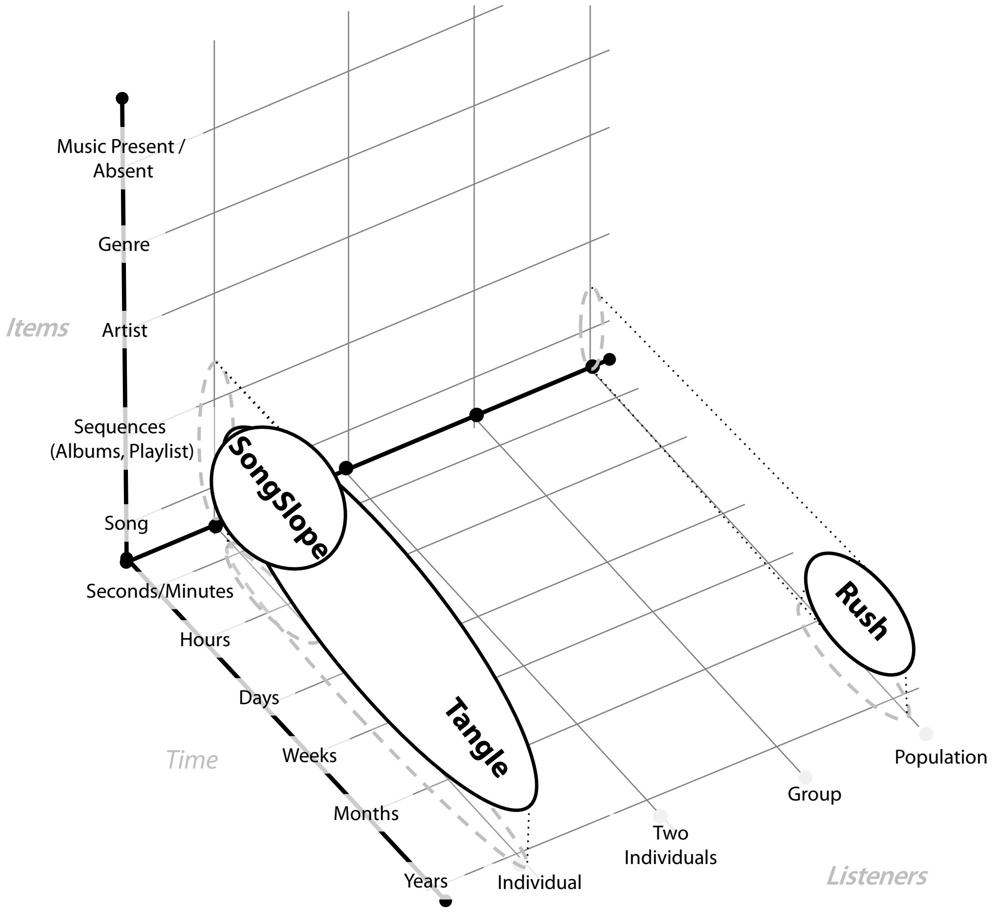

Figure 6.1 The three visualizations presented in this chapter in the listening history design space. Their focus lies on single songs and song sequences instead of time. Song relationships can also come from the whole last.fm population through collaborative filtering.

Tangle is a playful visualization, similar to *Strings*, that allows exploring the connections between songs in a personal listening history. All tracks are nodes in a graph and in-sequence events in the history form edges.
SongSlope finally is a visual plug-in for a media player that shows for the currently playing song its predecessors and successors. This lets the listener rediscover old songs and also create playlists on-the-fly by selecting an old listening session as a new playlist.

The casual approach of the visualizations in the last chapter is extended in this one: While the previous visualizations were all desktop-bound even if they were platform-independent, a PC screen showing a dedicated visualization tool is not necessarily the best way to create a playlist or listen to music. Music listening is increasingly moving towards more flexible and mobile setups and music tools should be able to help people without forcing them to their desktops. Therefore, the prototypes in this section are no longer necessarily meant for desktop use or present explicit analysis

tools. They are made for mobile devices (*Rush*) or integrated into existing media player software (*SongSlope*). They still rely on visualization techniques as the best way to visually represent characteristics of the data (connections between tracks). Figure 6.1 shows the tools' positions in the design space. What is immediately visible is that song relationship is mostly an all-or-nothing process regarding the *listeners* dimension. Either the combination of inherent and event-based relationships is taken into account for a single listener, or only the latter is used through techniques like collaborative filtering. Also, *time* is no longer in the focus and visualizations can concentrate either on short or long timespans. Aggregation and playcounts hold more meaning about the relationships.
Similar to the last chapter, I describe the design processes and motivations that went into each of the prototypes. I also present studies along the way and for summarizing their impact.

6.1 Rush

With an average song length of a few minutes, listening to music can be an involved activity: Choosing a new song every time the last one is through means that doing anything else in parallel is probably not possible. Usually, people rely on complete albums or playlists for listening to music. While listening to an album is convenient but means passing on other artists, creating a playlist can be an effortful process: Combing through one's complete music collection with thousands of songs to find the right ones can take even more time than switching songs manually in the first place (the peculiarities and theories on playlist or mix tape creation are described in section 2.2.5). In recent years, recommender systems have matured enough to be used as automatic playlist creators. Research on listener behavior (e.g., skipping [125] or genre taste [178]) and music content analysis (e.g., audio similarity [134, 50]) enabled the creation of tools that cut a swath through the collection. The underlying technology that powered automatic webradio such as that of last.fm or Pandora[1] became available in desktop media players with iTunes Genius or Microsoft's Smart DJ.
While these automatic playlist generators are helpful in avoiding the costly search through the whole collection, their results are often predictable and not necessarily appropriate. Even when the playlists are based on a person's taste (e.g., in Pandora) they might not fit for the current situation. Also, the simple voting mechanisms that are available (mostly binary like or dislike options) do not take into account a listener's history with a song: Based on the song's position in its lifecycle (cf. section 2.4.2) the listener might just have enough of it at the moment, but like it in general. Unfortunately, most playlist generators do not provide much in the way of customization: As the name 'Genius' suggests, playlist generation is a one-off process where the listener can either

[1] http://www.pandora.com

accept or discard the resulting playlist as a whole. The only way to influence the results is to choose a different *seed song* on which the playlist is based on.
All these aspects suggest the need for a more interactive way of playlist creation. In order to keep it manageable and useful, it should adhere to the general approach of automatic recommenders and provide some form of limitation for the number of songs, but not be completely inflexible. One reason for the restrictions of playlist generators is convenience: Solutions such as Satisfly [128] allow adjusting various musical aspects of the resulting playlists but can be overwhelming for regular listeners. Also, defining suitable constraints takes much more time than simple selecting a seed song and a length. Therefore, a possible solution not only has to be flexible but also easy and convenient to use.
Rush (see figure 6.3 and [10]) is based on the idea of replicating the manual listening process of selecting song after song but doing so in a more constrained fashion. After selecting a seed songs (similarly to automatic playlist generators), a set of five new songs is suggested that go with the original one. Once the listener selected one of them a new suitable set is suggested and so on. This allows flexibility not only for the length of the playlist (the process can be ended at any point), but also for each individual item to adapt it to the current situation and mood.

6.1.1 Design Process

Music listening and playlist building are increasingly shifting towards mobile use. As explained in chapter 2, music listening on-the-go is shaping up more to be the rule than the exception. Due to that, I wanted to make *Rush* available for mobile use. Deploying an application for mobile instead of desktop use brings with it a new set of requirements and constraints (see section 2.4.1). Possible interruptions and subsequent resuming have to be taken into account, and interaction is shaped by external influences (e.g., a shaking subway car). Also, interaction cycles are shorter and not necessarily the center of attention. I also decided to take a smartphone with touchscreen input as target device which brought with it its own set of problems regarding accuracy of input (c.f., "fat finger problem" [152]) and ergonomics. During the design process, I especially thought about this last part and how the device should be held and the interface be displayed. A main idea from early on was to make the interaction fluent: Usually, interaction processes are categorized through short bursts of interaction (targeting, mouse clicks) and long pauses. In *Rush*, I imagined the degree of interaction to be more constant: All control should be based on moving the thumb on the display and changing its position for navigating and selecting items. Initially, putting the thumb down should start the interaction and lifting it up would signal that the playlist is finished.
The underlying data structure of *Rush* is that of a graph (see figure 6.2). With the seed song as a root node, all other songs are connected based on their similarity. Songs can of course appear more than once within the graph, as they might share similarities with several tracks. For the first version of *Rush*, I relied on using last.fm's similarity data,

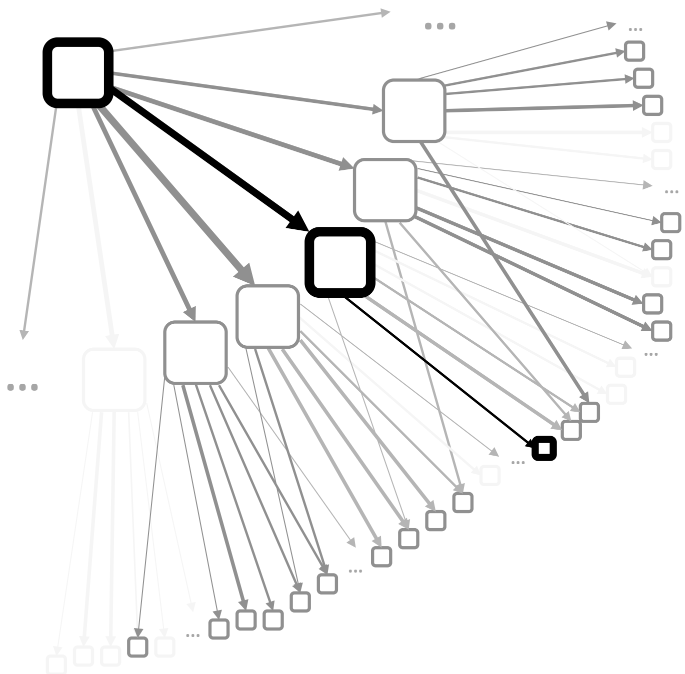

Figure 6.2 All songs are connected in a similarity graph. *Rush* presents five-song sets of this graph in sequence (dark grey arrows and boxes) and lets the listener select one song for each of these sets (black arrows and boxes).

which is based on collaborative filtering of all of their collected listening histories. It is possible to access a list of similar songs and confidence values for this similarity for each track via their API. Of course, building a personalized version of *Rush* with a custom similarity graph would be even better (see section 6.1.3).

Building a playlist with *Rush* works as follows: After manually selecting a seed song from a list of available tracks, this song and five related songs are shown. This represents a section of the similarity graph. As there are usually more than five candidates, some kind of selection algorithm has to be applied (see the evaluation section below for different approaches to that). Navigation and selection in *Rush* happen with the same continuous gesture: By putting a finger on the screen, all visible songs start sliding into the opposite direction (for example: putting the finger on the lower part of the screen pushes all songs upwards, while putting it on the upper part does the opposite). The further away from the center of the screen the finger is, the faster the sliding. Selecting a song for the playlist and accessing new suggestions happens with the same gesture: By crossing a song completely with the finger (i.e., drawing a line through it) it is added to the playlist and suggestions based on it are displayed. Multiple songs in the same row can be crossed in one stroke to add them all. Another crossing removes one or more songs from the playlist. Simply tapping a song also shows recommendations based on it without adding it to the playlist. Initially, lifting the finger was marking the end of

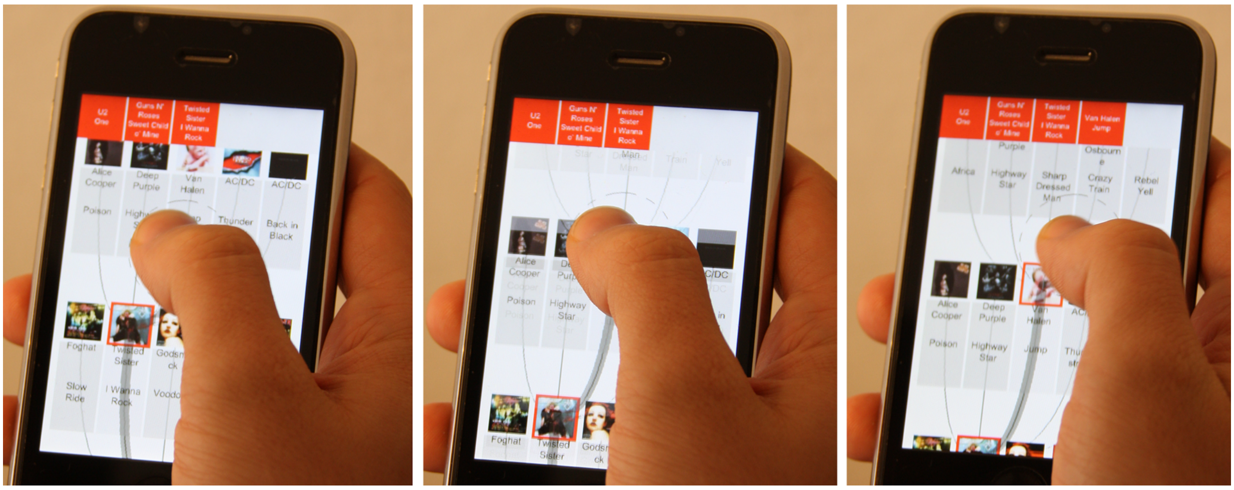

Figure 6.3 Creating playlists manually or automatically is either time-consuming or restricting. *Rush* is a mobile application for interactive playlist generation that allows quickly creating playlists suitable for the current situation (Source: [10]).

the playlist process, but early tests showed that especially in a mobile context having to press a finger to the screen at all times is awkward. Therefore, the playlist is finished by shaking the device, a gesture sufficiently different from the others to not be executed accidentally.
The resulting playlist is shown in the upper row of the screen (see figure 6.3) and can be navigated with swiping gestures.

Comparison with Dasher

Rush's approach to navigation and selection is similar to the *Dasher* interaction technique [171]. *Dasher* was developed for text entry on mobile devices and is based on a related principle: The system starts by displaying all 26 letters of the English alphabet. By clicking/touching somewhere on the screen, the underlying canvas starts moving towards this position. Letters are selected by crossing their outer boundaries. Once a letter is selected, the size of the next row of letters is adjusted based on probabilities. The system uses a corpus of English words to predict the next suitable letter in a row. After choosing the letters *'THEI'*, for example, the letter *'R'* is enlarged to make it easier to select. *Dasher* was very successful and had several incarnations in handheld devices, touch- or pen-based, and was also used for text entry via eye tracker [172] and brain-computer interfaces [176].
Rush is also based on the general principle of *Dasher* of combining navigation and selection in a single continuous gesture and gearing choices towards suitable ones. Differences are that *Dasher*'s approach uses a far higher number of items (26 letters versus 5 songs) and always provides the full set of available choices (which is impossible with a large music collection). It also frames the interaction differently: While *Dasher* talks about "[a] continuous zooming-in of the users' point of view"[171], *Rush* navigates along

the edges of the graph without the illusion of zoom. Also, *Dasher*'s text entry task has a clear goal of "typing" a single letter, while *Rush*'s goals are framed more vaguely: Songs are often good enough to fit, several songs from one row of suggestions can be added and discovery is always intermingled with the creation of the playlist. In *Rush*, the suggestions often trigger the wish to add a song to the playlist and not the other way round. Thus, the system also presents help for the uninspired, something which would be inappropriate for *Dasher*'s text entry use case.
While there are similarities between the two systems, one thing that has never been addressed in all the *Dasher* incarnations is the question of which direction the interaction should happen in. *Dasher* always simply flows to the right. Depending on how the mobile device is held and which finger is used for navigation, various combinations are imaginable. Occlusion is also a factor in this regard. To make sure that I would arrive at a suitable solution for *Rush*, I performed a formative evaluation that compared different approaches to this question for a simple navigation task.

Formative Evaluation

Mobile device are usually rectangular, which means that they can be held comfortably either horizontally or vertically. Subsequent interaction happens either with two thumbs or the index finger of the non-holding hand. Depending on the use case, an application can favor either of these *device orientations*. *Rush* and *Dasher* both have additional parameters: The *interface orientation* determines whether new items appear to the left and right (*horizontal interface orientation*) or above and below (*vertical interface orientation*) the currently active item. The closely connected *interface direction* controls if the "forward"-direction where new items appear lies to the left or right (in the *horizontal interface orientation* case) or above or below (*vertical interface orientation*). Finally, when interacting with the index finger, either that of the dominant or non-dominant hand can be used (*used hand*).
When designing *Rush*, I performed a formative evaluation first, to determine which of these settings would work best for the application (cf. [10]). The task the participants had to perform was to select ten digits using the *Rush* interface (see figure 6.4). Their finger's position relative to the center of the screen determined movement of the underlying canvas and the distance to the center the speed. I used an Apple iPhone 3G for the study with a 320x480 pixel touch-screen, so I capped the speed at 160 pixels distance to prevent advantages for vertical orientations. The sequence of digits was predefined and randomized and shown at the top or left end of the device (depending on the orientation). Participants received five suggestions for each subsequent digit and had to select the right one as fast as possible. When making a mistake (see figure 6.4 right) they had to deselect the wrong one first before continuing. I counted the time for each task (measured from putting the finger down to selecting the last digit) and the number of errors (either selecting a wrong digit or deselecting a right one).

The study had a within-subjects design. Each participant had to perform the task for a combination of each of the variables (the order was randomized): *Device Orien-*

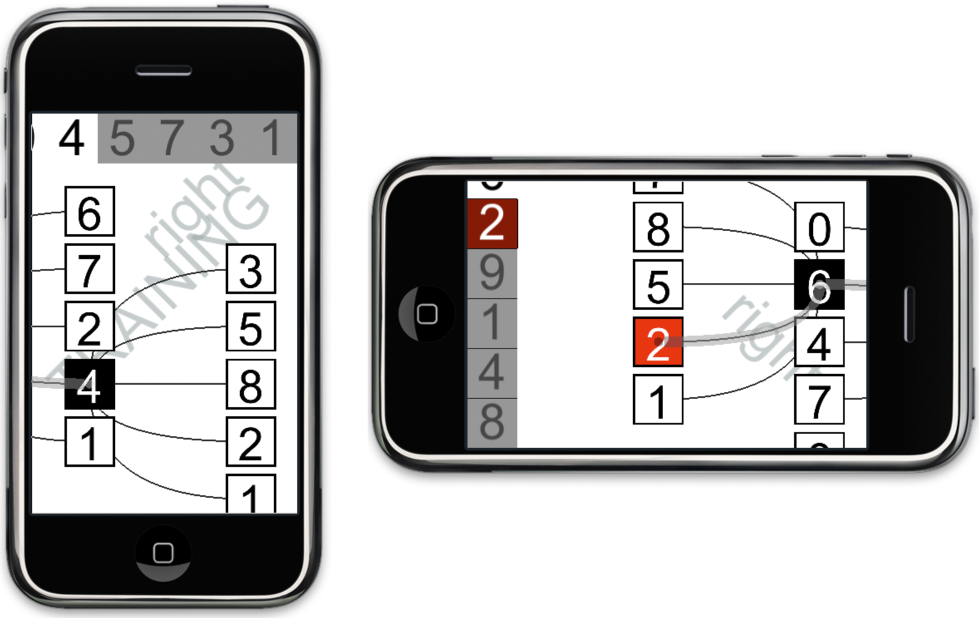

Figure 6.4 In the formative evaluation for *Rush*, participants had to select ten numbers with different interface settings. Both examples show horizontal interface orientations, but the left side has a vertical device orientation and right as the interface direction, while the right side shows a horizontal device orientation and left as the interface direction (Source: [10]).

tation (horizontal, vertical) x Interface Orientation (horizontal, vertical) x Interface Direction (up-right, down-left) x Used Hand (left, right) which resulted in a total of 16 rounds for each participant. Note that I merged the *up* and *right* and *down* and *left* cases for *Interface Direction*, as they were dependent on *Interface Orientation*. Before performing in each of the 16 timed runs, participants always tried the type of interface in a training run, which resulted in another 16 performances (which were not timed).

Participants and hypotheses For the formative study, I recruited 12 participants from our university (21 to 32 years, avg. 27.4, three female, two left-handed). Each of them had at least some experience with handheld touch-screen devices. In addition to learning what the fastest settings for orientations and direction were, I also wanted to learn more about the influence of occlusion and handedness. I had three hypotheses (cf. [10]):

H1: Occlusion leads to higher task times.

H2: The dominant *used hand* would fare better than the non-dominant one regarding time and errors.

H3: Identical *interface* and *device orientations* would improve task times.

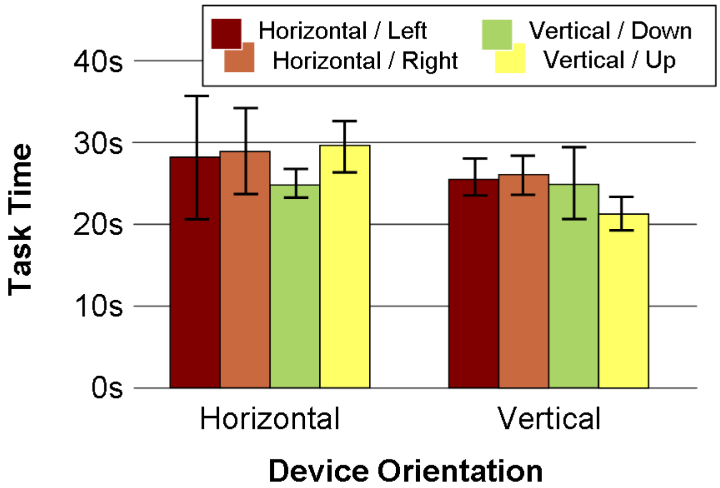

Figure 6.5 The results of the formative evaluation showed that the vertical *interface* and *device orientations* were significantly faster than the horizontal versions (Source: [10]).

Results To analyze the effects of the independent variables on task times and error rates, I performed a repeated measures ANOVA on the results. Additional interaction effects between the variables were examined using subsets of the data. "All post hoc pairwise comparisons used Bonferroni corrected confidence intervals for comparisons against $\alpha = 0.05$" [10].
In a preliminary test I looked for indications for **H2** and the effect of handedness on the results. During performing the study, participants repeatedly complained about having to use their non-dominant hands, but I wondered whether this aversion would also show in the results. For the two left-handers the results showed an average task time of 24.12 seconds for the left and 26.46 seconds for the right hand. Yet, there was practically no difference for the right-handed participants (26.34 seconds for the right and 26.39 seconds for the left hand). The small number of samples for left-handers and the statistical analysis that found no effect between handedness and either of the dependent variables, led me to reject **H2** and remove it from the subsequent analyses.
Significant results were found, however, for the design variables: "We found significant main effects on completion time for both Device Orientation ($F1,10 = 13.056$, $p < 0.005$) and Interface Orientation ($F1,10 = 7.094$, $p < 0.024$). There were no significant interaction effects in our data. Overall, the Vertical Device Orientation (M=24.49, SD=1.29) was faster than the Horizontal one (M=27.85, SD=1.28). The Vertical Interface Orientation (M=25.16, SD=1.23) was also faster than the Horizontal Orientation (M=27.17, SD=1.28)" [10]. In concurrence with **H3**, the combination of *vertical Device Orientation* with *vertical Interface Orientation* proved to be the fastest one (M=23.13, SD=1.76) and was on average 4.05 seconds faster than all the other combinations.
There was no significant effect between *Interface Direction* and task times, but for a combination of vertical *Device-* and *Interface Orientations* an upwards movement was fastest (M=21.23, SD=1.16). The relative speed of vertical device and interface orientations and

the upward direction have little occlusion compared to other combinations (for example, horizontal interface and device orientations, direction right and used hand also right), which speaks for confirming **H1**.
Regarding the error rates, there was a significant effect "for *Device Orientation* ($F1,10 = 6.139$, $p < 0.033$) but no significant interaction effects. Post-hoc multiple means comparisons revealed that the Vertical Device Orientation (M=0.53, SD=0.15) performs better than the Horizontal one (M=0.86, SD=0.14)" [10].

Influence on the design

The formative evaluation brought some clarity for the design of the final version of *Rush*. While I had always favored a one-handed, thumb-based use of the interface, the results had shown that this combination of device and interface orientation produced the best task times. Therefore, I also used this combination together with an upwards movement direction for the final implementation.
One thing that I also found during the evaluation was the importance of occlusion not by the hands but the display boundaries: When touching an item, the new suggestions appeared beyond the display boundaries. This usually caused participants to move the finger as far towards the boundaries as possible (to increase the movement speed), but also often led to an erroneous selection once the items appeared. A related problem was caused by movements orthogonal to the interface direction: If, for example, the horizontal interface moved from left to right, participants were still able to move up- or downwards, which often required them to first move vertically before being able to see all five suggestions. This caused some frustration.
In the final version of *Rush* I therefore restricted movement to the interface direction axis and also elongated the items to prevent erroneous selections. In the formative version, the movement of the canvas was also determined by the finger's position relative towards the actual center of the screen. In the final version, I shifted this virtual center towards the bottom of the screen, to leave more space for new items to appear and making decisions easier. Also, sometimes the five suggestions for the last song touched contain no suitable track. By touching another song in the same row these suggestions are replaced (instead of displaying the combined ten suggestions, for example). If this replacement happens while one of the suggested songs has already been added to the playlist, only the four remaining songs are switched.

6.1.2 Evaluation

While the results of the formative evaluation allowed me to adjust layout attributes of the interface suitably, I had not tested all possible variables. I had arrived at the number of five suggested items, for example, through the available screen space, that allowed displaying five items next to each other without sacrificing readability. But the question remained how to select these five and in which order to present them. Additionally, I

wanted to test whether *Rush* worked for its intended purpose of building a playlist.
One important aspect is the order in which songs are placed on the screen as this placement can be used to encode information. Instead of sorting the songs alphabetically or placing the most appropriate suggestion at the left- or right-most position, I decided to use the middle slot for that. Based on interaction with the thumb the middle slot was the ergonomically most appropriate position. This placement logic also made it possible to recreate the results of an automatic recommender system simply by drawing a straight line through all the middle items.
Suggestions in *Rush* are based on last.fm similarity values for songs. This makes sure that the suggestions have at least some relationship to the original item and help in building a consistent playlist. Similar songs come with a confidence value. It is therefore possible to sort suggested items along their similarity. Just taking the five most similar items might, however, not be the best strategy, as some kind of variation is usually desired. On the other hand, randomly picking songs from the similar items can yield either good or bad results. I wanted to test these different *suggestion strategies* in another study. I compared three strategies:

Top-5: This strategy simply took the top-five most similar songs and placed them from the center of the screen outwards. The most similar song was thus in the center and the least (most) similar ones in the periphery.

Random: In this strategy, the system picked five songs at random from the list of similar songs and placed them in an equally random order. The middle item was thus not necessarily the most appropriate suggestion.

Hybrid: For this strategy, I wanted to combine the appropriate suggestions of **Top-5** with the lack of a lock-in effect. The middle item is the most similar song, while the two next to it show suggestions from the middle of the similarity list. The two items in the periphery are the least similar songs from the list and thus allow breaking out of this musical direction without breaking the consistency.

In the study, I not only wanted to compare these three different strategies but also see how well *Rush*'s results would fare against automatic recommender systems and manual approaches. This comparison should only happen quality-wise based on the resulting playlists. That the three approaches were very different in required time and flexibility was obvious beforehand.

Song set

The study setup required a music collection based on which the playlists were created. To keep the results comparable, I created a predefined set of 3,900 songs that was used for every participant. The goals for this collection were:

- to only use relatively well-known songs to make sure that participants knew most of them

- have a varied collection so that every participant could find something they liked without frustrating them
- make sure that the lists of similar items for each song contained at least ten entries to songs of other artists, so the suggestion strategies would yield different results

To create the collection (cf. [10]), I first collected all-time favorites for the genres Rock, Pop, and R'n'B on the web. I then wrote a python script that downloaded similarity lists for each of these songs from last.fm and checked how many listeners each of these similar items had. If there were at least 500,000 scrobbles for a song it was deemed sufficiently well-known and added to the collection. Also, each similarity list had to contain at least ten other songs by different artists to be viable. Once the song set had reached the size of a common music collection I quit the script and started another one that downloaded album covers and thirty-second samples for each song from Amazon.com. Audio samples would make it easier for participants to gauge the appropriateness of a song in a playlist. They were played while touching an item and contained not the beginning of the song but an identifiable section in the middle (mostly the chorus).

Study design

This second study again had a within-subjects design. Each participant was asked to select a seed song and create four ten-song playlists based on it, three using *Rush* and one manually. Each incarnation of *Rush* followed a different suggestion strategy (*Top-5*, *Random*, *Hybrid*) and these strategies were counterbalanced among the participants. Additionally, a set of three playlists was created automatically based on the seed song without the participant's involvement.
In the manual condition, the participants had access to all 3,900 songs of the whole collection including album art and audio samples and were also able to see the full similarity lists. The automatic playlist generator worked internally similar to *Rush*: It had the same three suggestion strategies and replicated interaction by picking one of the five suggested songs at random. This resulted in another three playlists. In the end, I arrived at seven playlists for each participant (three *Rush*, three automatically, one manually).
Besides building playlists, participants were asked after each *Rush* condition how they liked the suggestions and how random they appeared. In closing, they also rated the system on a modified IBM usability satisfaction questionnaire [104] and provided demographic information. Finally, they were asked to rank the seven playlists.
For this second study, we had another 12 participants (4 female, 2 left-handed, 24 to 35 years, avg. 28 years) and four of them had participated in the formative evaluation. They again stated experience with touch-based devices.
We also had another three hypotheses about the results:

H1: Creating a playlist with *Rush* would be faster than with the manual tool.

H2: *Rush* playlists would be ranked higher than those created automatically.

H3: *Rush Hybrid* playlists would be better and the suggestions would be preferred to *Rush Random*.

Results

When looking at the times participants took to create playlists, the *Rush* conditions fared better than the manual one, as expected in **H1**: *Rush Top-5* took on average 123.8 seconds, while *Rush Hybrid* was finished after 142.1 and *Rush Random* after 162.3 seconds (cf. [10]). Taking the participants' comments into account, they often just gave up on the *Top-5* condition (one participant commented that "There are always the same songs!") which is one explanation why it was fastest. The manual condition took on average 388.6 seconds, thus confirming **H1**.
On the qualitative side, participants in general liked *Rush*'s usability and were happy with the number of five suggestions (58% found it neither too high nor too low). The only sub-par question was on the operation speed that resulted in an average ranking of 2 on a five-point Likert-scale, with 1 meaning "too slow". Participants had also been asked to rank the tools they had available and favored the the *Manual* condition (3 wins using the Condorcet Ranked-Pairs system). Follow-ups were *Rush Hybrid* (2 wins, 1 loss), *Rush Random* (1 win, 2 losses), and *Rush Top-5* (3 losses).
Measuring the quality of playlists is difficult in general (cf. [5]). While I had asked participants to rank the seven playlists resulting from their study run, it was clear that they would be biased towards their own creations. Ranking their own playlists was insofar suitable as they represent a personal look at music and could only be rightfully evaluated by their creators. On the other hand, these creators had also spent several minutes on each of these playlists (even longer in the manual condition), so it would put them into a state of cognitive dissonance when voting for the automatic playlists and thus devaluing their own work. Therefore, I additionally made the resulting playlists available online and asked anonymous reviewers to rank the playlists again. I received 10 rankings in total and used them as a more impartial evaluation of the playlists.
To analyze the participants' playlist rankings I again used the Condorcet Ranked-Pairs system. As expected, *manual* playlists won (6 wins), "followed by *Rush Hybrid* (5 wins, 1 loss), *Rush Random* (4 wins, 2 losses), *Automatic Random* (3 wins, 3 losses), *Rush Top 5* (2 wins, 4 losses), *Automatic Top 5* (1 win, 5 losses), and *Automatic Hybrid* (6 losses)" [10]. The more independent online vote painted a different picture: "*Rush Hybrid*, *Rush Random* and *Automatic Hybrid* are tied for first place (4 wins, 2 unresolved), followed by *Automatic Random* (3 wins, 3 losses), *Manual* (2 wins, 4 losses), *Rush Top 5* (1 win, 5 losses), and at the last position *Automatic Top 5* (6 losses)" [ibid].
Remarkable about these results is especially the drop of the *manual* playlist from the first to the fourth place and the subsequent rise of *Rush Hybrid* and *Rush Random*. For one, this shows the assumed participants' bias towards their own creations. It also means, however, that the manual, freely-defined playlists had a lower objective quality than the ones created within *Rush*'s restrictions. Therefore, the five item suggestion of *Rush* could even help in creating better playlists.
Regarding the hypotheses **H2** and **H3**: Both sets of rankings had *Rush Hybrid* and *Rush*

Random above the automatic results, confirming **H2** for these conditions. Also, *Rush Hybrid* came out ahead close before *Rush Random* in one of the playlist rankings (they were tied in the second one) and in the tool ranking. These differences were so close, however, that they are only barely confirming **H3**. The much more elaborate *Hybrid* suggestion strategy thus was not much better than simply picking random similar songs.

6.1.3 Integrating personalization

While *Rush* has proven to be a suitable interface for creating personalized playlists quickly, my prototypes were only based on aggregated versions of listening histories. All of last.fm's similarity values came about through collaborative filtering of millions of histories, but the results only provide an objective version of song relationships. It would be much more interesting (and probably satisfying) to work with a personalized version of *Rush*.

Integrating personal listening histories into the system would bring several advantages compared to relying on objective similarities: First, using the listening history to extract similarities based on the distances between songs instead of the aggregated overall versions could increase the satisfaction. While the collaborative filtering results are usually appropriate, every listener has her or his own little quirks and stories behind songs (see section 2.4.1) and they do not necessarily have to conform to the overall perception of music. These peculiarities are usually hidden in the aggregated version.

Second, the system could also use the data to include personal playlists into the suggestions. As a listening history does not only contain single songs but also repeating sequences, they can also be used to improve the resulting playlists. One option would be to fill the default middle item with a popular (i.e., frequently repeated) song sequence from the listener's history and allow either just listening to it by always selecting the middle item or deviating from it with one of the other four suggestions.

Finally, a listening history allows the system to make more informed decisions about which songs to display. The five available suggestion slots in *Rush* are not a lot and each "wrong" suggestion restricts the listener more in her or his playlist. Such wrong suggestions can come in different forms. For one, the *Rush* principle presumes that the listener knows all of the suggestions to allow a quick selection without first listening to the songs. A (meticulous) listening history contains exactly this information and can help in avoiding unknown songs. Also, each song is at some position in its lifecycle (cf. section 2.4.2) and might just now be not interesting for the listener due to overexposure (even though its playcount is high). When knowing the exact listening history, the system can also avoid such songs.

All in all, *Rush* is a promising interaction technique for creating playlists: It presents a middle ground between purely automatic and patronizing playlist generators and flexible but tedious manual creation. Additionally, it is made for mobile devices, where a large part of music listening anyway happens nowadays. It shows a way how listening histories can be valuable assets for easing future music listening.

6.2 Tangle

Strings (see previous chapter) was a first attempt to understand the patterns available in a listening history. Its focus was thereby on the *time* dimension, and interpreting the data as part of a personal life story. The appearance of song instances throughout the timeline made it difficult to see the *item* patterns and the connections among the songs themselves.
Tangle (see figure 6.6 and [12]) is a related project that also relies on simple access and minimal interaction, but puts its emphasis on the song-to-song relationship. The initial motivation for *Tangle* was that I wanted to have a tool that allowed me to see how closely knit the relationships between songs in a listening history really were. Would repetitions only seldomly appear or repeatedly? How radical would the differences between two histories be?

6.2.1 Population and goals

The target group for *Tangle* is similar to that of *Strings*, but with a different focus. Therefore, this group of analysts would again have a highly diverse background with no guarantee for musical or visualization knowledge. Only knowledge about the listening history would be at hand and even though *Tangle* does not include temporal information into the visualization it is nevertheless aimed at people working with their own history. External analysts are not (directly) supported, even though they are also able to draw conclusions from the visualization (sans the explanations).
The actual goals for the analysis cover three different aspects:

A: Find patterns along the *item* dimension. This first and main goal establishes the focus of the visualization. Patterns within the *items* are not only popular relationships between songs, but also the general importance of a single song. So, if a song appears multiple times throughout the history, this should be just as visible as if it had been listened to only once.

B: Find sequence patterns within the data. Goal A focuses on single items, but looking at song sequences within the data (playlists, albums) might also bring interesting results. A second goal is therefore to make this sequential information available. Note that uncovering sequences does not necessarily require adding temporal information. Song sequences are, so to speak, on an ordinal and not an interval scale.

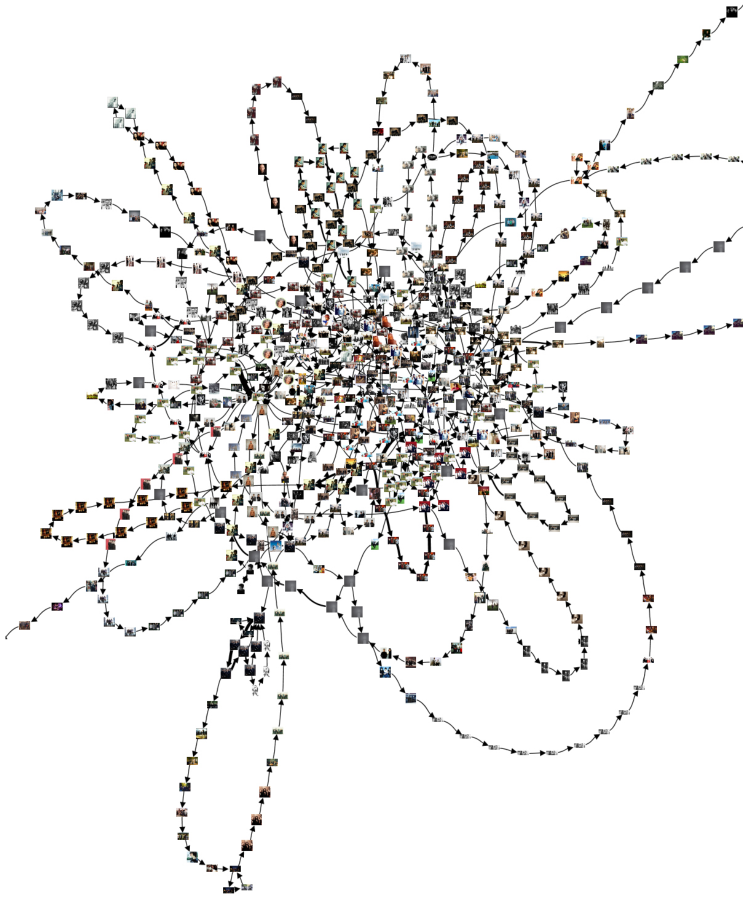

Figure 6.6 *Tangle* is a node-link diagram visualization for song-to-song relationships. Each node represents a song, while each edge stands for a sequential listening for the two songs.

C: Have easy-to-learn interaction and an immediate benefit from the visualization. This goal is identical to the respective goal in *Strings*: The application should provide simple interaction as it is not used often enough to warrant a long learning phase. The interface should be minimal and functionality easy to discover. Immediate benefits are always useful to receive relevant information more quickly and to inspire analysts to work with the visualization longer.

6.2.2 Design

The main idea of *Tangle* goes back to the problem of the time-line based visualizations (cf. last chapter). Each song has potentially multiple instances that appear throughout the timeline. Identifying the importance of a song is difficult and finding patterns beyond single songs (i.e., repeated song sequences) practically impossible. All visualizations of the last chapter had an (usually interactive) work-around for this problem, as instances of single songs were connected with arcs. In the case of *LastHistory*, these arcs also took sequences into account and highlighted repetitions of multiple in-order items. In *Tangle* I decided to solve this problem in a different way. Instead of connecting related instances, the unique song itself would be the main item of the visualization. Each song had only one representation. Listening to music was nothing more than visiting each of these songs one after the other or even repeatedly. Therefore, I opted for a simple node-link-diagram to display the data. Each unique song is represented by a node (displaying an image of the artist) and edges encode the listening: A directed edge connecting song A and song B means that song A has been heard first, followed by song B. To make it easier to see important sequences in the graph, each edge that has been crossed more than once (as the song-combination has been heard multiple times) becomes thicker (see figure 6.7). Popular sequences of songs therefore have strong, easily visible arrows connecting them and are thus better to see at first glance. In a way, this node-link-diagram is a timeline, spanning from the first to the last song of the history, that has been wound around each contained song.

Node-link-diagram commonly suffer from the problem of overlap. Especially if nodes have multiple edges, avoiding overlap and the so-called "hairball" in the layout is usually not possible. I chose a force-directed layout algorithm for creating the positioning of the nodes. Nodes act repellent on each other while edges try to contract as far as possible. The layout algorithm can be run continuously which also allows exploring the graph by interactively dragging nodes and following the changes in layout.

Interaction was kept as minimal as possible. I had panning enabled through dragging the background with the left mouse button and zooming via the mouse wheel. Hovering over a song shows a small pop-up with name of the artist and title of the song (to lower the visual clutter of the application this information is only shown on-demand). Finally, songs can be dragged around with the left mouse button and the force-directed layout ensures that related songs follow.

These design decisions support the three main goals in different ways. Having songs as the central aspect of the visualization helps both in seeing patterns for simple songs and sequences: Popular songs are automatically drawn towards the center of the hairball, while less popular ones reside on the fringes. Regarding goal B, the thickness of the arrows that connect songs makes the popularity of sequences again visible at a glance. The force-directed layout also causes interesting sequences (e.g., cycles, cf. 6.7) to become visible. Finally, the simple drag-and-drop-based interaction requires no learning and discovering that songs can be dragged can even happen accidentally when trying to pan the canvas.

6.2.3 Discussion

In a way, *Tangle* represents a social network analysis for music. When each song enters a form of relationship with another by being listened to in sequence, then *Tangle* uncovers the degrees of separation between the different tracks. Groups can arise via overarching

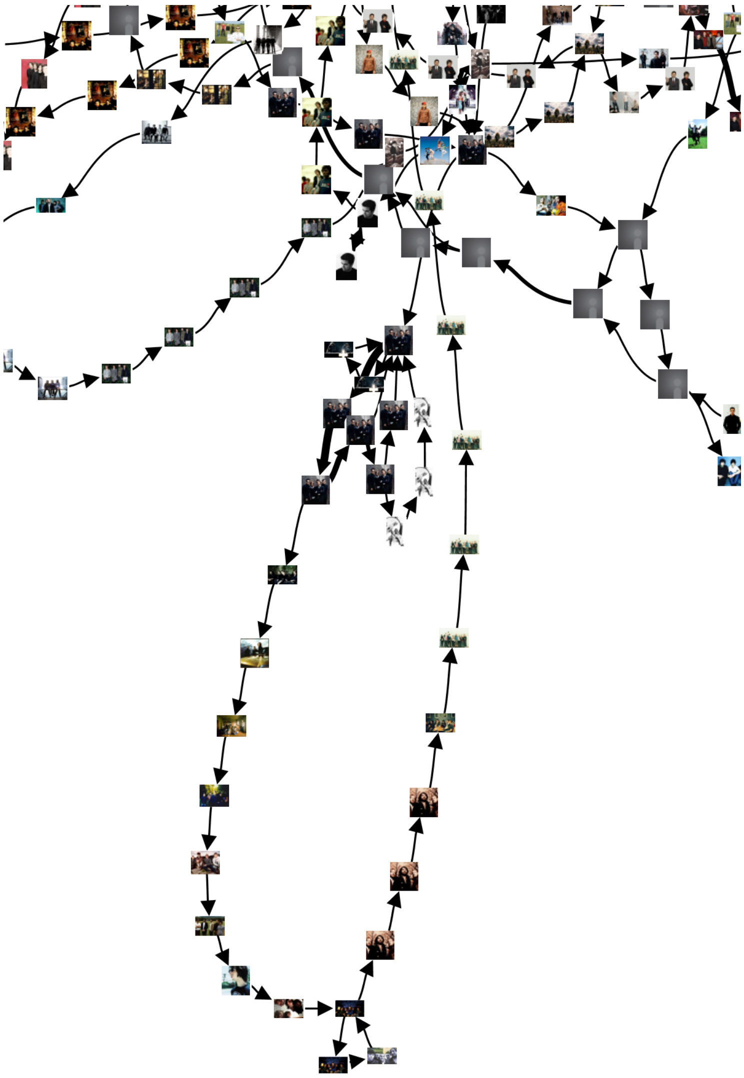

Figure 6.7 The number of repetitions of a two-song sequence is encoded with the thickness of the connecting arrows. Less popular songs wander towards the periphery.

commonalities of songs such as having been performed by the same artist or belonging to the same genre. Similarly, song sequences are stronger chains within such a music network and each repetition by the listener strengthens this chain. When relying on the song-to-song relationship a completely new perspective on a listening history unfolds. *Tangle* was, just as *Strings*, more or less a proof-of-concept, so I was not interested in its usability or other aspects that could have been uncovered through a summative study. It was aimed at the lowest common denominator for background knowledge and also only displayed artist and track names. It works in a playful manner, without major possibilities for interaction and analysis, even though certain patterns of the data can be seen. The minimalism for interaction also means that every benefit that can be gleaned from the visualization is immediately there. Dragging songs around can help in gauging

the importance of single items but does not contribute much to the overall picture.
What I also learned from *Tangle* was, how complicated even small music networks can become. The system was also implemented using the *prefuse* framework, and the costly force-directed layout algorithm meant that no more than 1,000 songs could be shown at a time. But even then, the entanglement was for a regular listening history so dense that a faster machine would not have brought more insight. Integrating abstraction mechanisms such as artists or genres could have helped here. And we tried another alternative, namely a "query-by-song" approach, in the next prototype, *SongSlope*.

6.3 SongSlope

Listening to music and analyzing music listening are two completely different activities. As the previous prototypes have shown, creating and using dedicated applications for either of them appears to be the best solution. However, especially the personal uses of a music listening history, namely reminiscing and rediscovery of old songs, could also be convenient outside of an explicit analysis context. Just like stumbling upon an old photo album or souvenir, accidentally hearing an old song (e.g., while listening to one's whole music collection in shuffle mode) can bring back memories and nostalgia. It might also trigger an inquiry about when this song has been heard last and in what circumstances and what other songs were relevant in that time period.
SongSlope (see figure 6.8) is a plug-in for the open-source media player*Songbird*[2] that supports such tasks. When we[3] developed it, our main goal was to bring casual analysis capabilities to situations where people listened to music. While some of the previous prototypes such as *LastHistory* also allowed listening to the music that was shown, their approach was more from the analysis side with playback capabilities tacked on. Launching such analysis tools was an explicit activity that happened at best every few weeks. *SongSlope*'s goal was to integrate the convenient reminiscing facilities that the analysis prototypes brought into everyday music listening.
To do this, we focused again on the song-to-song-relationship (cf. also figure 5.6). Each song is part of one or more sequences of songs. Each of these sequences is in turn part of one or more listening sessions. This hierarchy was the one basic idea for *SongSlope*. The other one came from the way people listen to music, namely song-by-song. Therefore, the application relies on the currently playing song for displaying additional reminiscing information about it.
Finally, we also wanted to do an online evaluation of the tool and see if it works as planned with any computer and last.fm listening history and if people are able to make use of it. During this evaluation, we were interested not only in general feedback about usability and usefulness but also learn more about the direct interaction with the tool by logging low-level interaction.

[2] http://getsongbird.com/

[3] *SongSlope* was part of Hans-Peter Dietz's Bachelor Thesis [45]

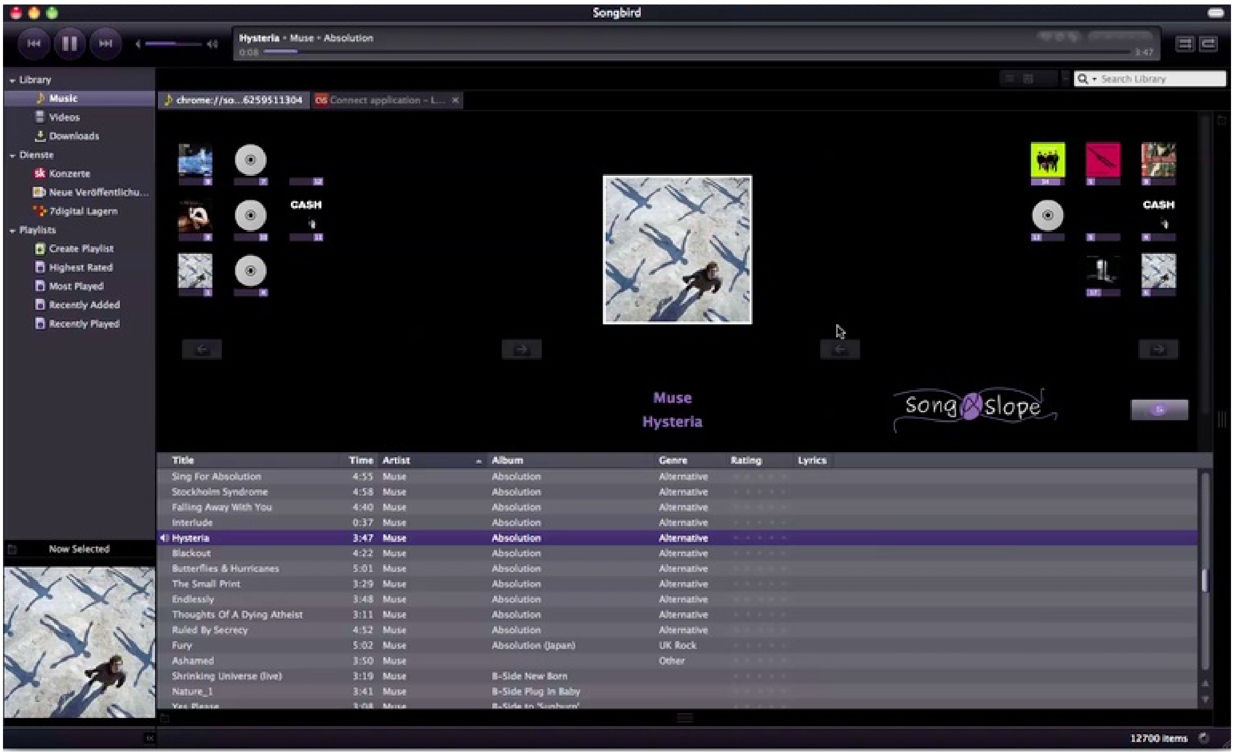

Figure 6.8 *SongSlope* is a visual plug-in for the open-source media player *Songbird*. It shows the currently playing song's predecessors and followers (Source: [45]).

6.3.1 Population and goals

The intended user population for *SongSlope* is slightly different than that of the other history visualizations. While a profile on last.fm is again a requirement to make use of the application, the lifelogging aspect is not nearly as prominent. An intention to learn about one's life, analyze patterns, etc. is not required and should not be supported by the tool. More sophisticated ways to access or filter the data should not be provided to keep the interface as clear and easy-to-use as possible. Therefore, *SongSlope* is entirely focussed on the song-to-song aspect of a listening history and works as an enticement to learn more about one's behavior. People who shun last.fm's data as lifelogging instruments might nevertheless enjoy getting access to this data in ways they can find useful: as predefined playlists or for rediscovery.

Lifelogging experts who want to draw something from their listening history can use the tool as an unobtrusive reminder and memory extension: When listening to a song they should be able to see what other songs might come to mind and what periods of time they had been heard. This becomes possible even outside of an explicit analysis context and simply by using a media player for listening to music.

When using *SongSlope*, the listener has the following goals:

A: Learn about song relationships in a known context. The main goal a listener has with *SongSlope* is to unobtrusively learn about their listening habits without having to do anything different than they would have done anyway. Reminiscing and rediscovering with the tool should be a background activity. Just as music playback was tacked on in *LastHistory*, analysis tasks should be tacked on with *SongSlope*. Accordingly, the interface has to be directly integrated with a music listening tool and present the visualization only as an alternative view on the current song.

B: Listen to music first. Actually listening to music must come first before all analysis activities. Therefore, the integration of the visualization with the media player should have no detrimental effects on the actual music listening. Also, listening to the displayed songs should be possible throughout the visualization, just as other convenience functions (e.g., adding a whole set of songs to the current playlist).

C: Allow access with easy-to-learn and minimal interaction. Similar to the previous prototypes, interaction and a corresponding interface should not make up the foreground of the application. The interface should be self-explanatory and simple, with the paradigms of immediate benefit and a step-wise discovery of additional features in place.

These goals should be supported in a regular media player and the interaction should be as conveniently as possible.

6.3.2 Design

The main goal of having reminiscing activities in a common usage context meant that including *SongSlope* into an existing media player would probably work best. Most of them (e.g., Apple iTunes, Winamp, Amarok) support plug-ins from third-party developers. In addition, these media players often present add-ons in dedicated catalogs which would make being discovered far easier and lead to a larger numbers of downloads. After examining the available target platforms, we decided to implement *SongSlope* for the open-source media player *Songbird*. *Songbird* has many advantages compared to other media players: It is cross-platform (Windows, Mac, Linux), explicitly extendable with third-party visual plugins (so-called *Media Views*), supports add-ons written in XUL[4] (an XML language based on web-standards) instead of a proprietary environment (e.g., iTunes), allows access to the full music library and metadata, and has a very strong on-board support for last.fm functionality (last.fm radio and scrobbling are integrated with the default installation). Finally, as we mainly learned during the implementation phase, the developers of *Songbird* are very invested in their product and provide quick and extensive support for add-on developers.
Songbird's strong integration of last.fm in combination with the full access to the media

[4] https://developer.mozilla.org/En/XUL

library meant that we could make *SongSlope* very convenient to use. Configuration happens completely automatically without any involvement of the listener. If the last.fm add-on is active and connected, *SongSlope* downloads the listening history and creates the necessary connections with the available songs in the library in its own database. Once this first setup has been completed it runs as its own *Media View* and can be switched on as an alternative to *Songbird*'s other views (e.g., a cover flow clone or song lyrics). The listening history is updated in regular intervals when restarting the application. *Songbird* also automatically scrobbles every song heard and adds it to the listening history.

SongSlope's visual interface is trimmed down to the minimum (see figure 6.9). The main

Figure 6.9 In *SongSlope*, the song surroundings of the currently playing track are shown. Songs that came before it are on the left, while songs that came after are on the right. The small bars denote the length of the respective listening sessions (Source: [45]).

element is the currently playing song in the middle. To its left, songs that came before it in other listening sessions are shown with their album art. Similarly, following songs are displayed on the right. All icons are sorted by the date of the corresponding listening sessions. By hovering above one of the song icons, additional information about artist and title is shown. A little bar below each icon shows the length of the session where this combination appeared. *SongSlope* works with listening sessions, similarly to *Strings*, and interprets an hour of inactivity as the start of a new listening session. Dividing the history into chunks makes it easier manageable. Also, one can argue if a song-to-song-relationship still exists between songs in different listening sessions. One peculiarity of *SongSlope*'s song display is that songs that appear several times in combination with the current song also appear several times to its left and/or right. This allows seeing important combinations without introducing more visual complexity through items with different sizes.

One interaction consistent throughout the tool is that clicking on a song starts its playback. This is visually supported by an animation that translates the item currently in the center towards the left (as it becomes a predecessor of the new song). Clicking on one of the session bars below an item changes the visualization and puts its focus on the respective session.

This session view (see figure 6.10) shows all songs from that session, together with

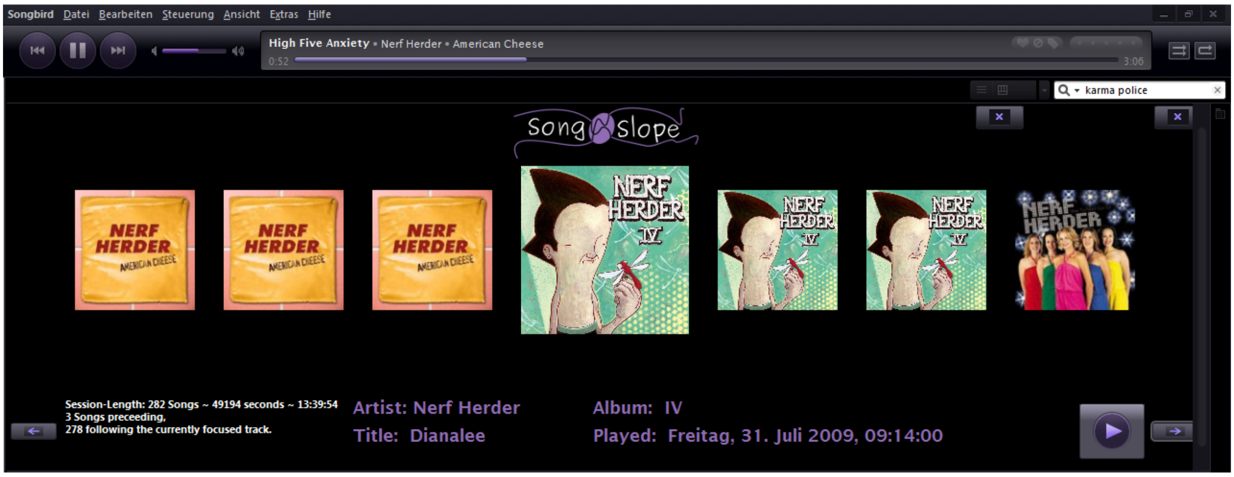

Figure 6.10 By clicking on one of the predecessors or followers, the respective listening session is shown and can be explored. All songs can also be added as a new playlist (Source: [45]).

metadata such as the date, number of songs and length. Clicking on the arrows beneath the song in focus shifts the display to the left or right to see the other songs of the session (the available screen space only holds seven songs at a time). The lower right corner of the visualization shows a play-button that allows using all songs of the session as a new (old) playlist.
These characteristics make *SongSlope* an unintrusive addition to the *Songbird* player. The listener can switch it on at any time and see personally related songs to the current one. By enlarging the player window while doing something away from the computer, the song surroundings are available at any time.

6.3.3 Reception and evaluation

SongSlope was made available in the *Songbird* add-on catalogue[5] in August of 2010. We added an introductory video and some explanatory text which was enough for understanding the reduced interaction of the tool. After a month it had been downloaded around 1,500 times (at the time of writing and 13 months later downloads are at 12,500).
The evaluation approach for *SongSlope* was two-fold: On the one hand we provided an online questionnaire and a pop-up reminder after ten days of use that asked people to fill it out. Similar to *LastHistory*, the pop-up disappeared once the link had been clicked. The survey contained the usual questions about demographics, usability and free form questions about gained insights.
The second approach was low-level interaction logging that contained detailed information about clicks on various parts of the interface. *SongSlope* collected this information

[5] http://addons.songbirdnest.com/addon/1902

and sent it automatically to our server after launching. While this type of feedback was more convenient for the listener it of course had more problems regarding privacy. We therefore made sure to restrict the collected data to aspects that could not be used for identifying people (e.g., we did not collect song titles heard) and asked the listener when first launching the media view if the data could be collected and informed her or him about its specific nature and what we would do with it. We hoped to gain more information about actual usage of the tool and be able to identify problems easier with this type of data.

Results of the questionnaire

Of the 1,500 downloads, only 30 people started filling out the questionnaire and of those, 16 completed it. Regarding the demographics, we had a predominantly male audience (87.5% to 12.5%) with an average age of 23.75 years (ranging from 13 to 46). Almost half of them had an academic background (43.75%) while most of the rest did not provide occupational information.
Asked about the way they use *SongSlope*, participants stated mainly for *rediscovering my music* (31.6%), followed by *analysing my listening history* (29%) and *generating playlists* (21.1%). They also spent about 1.5 as much time in the main view of the application as in the session views. Surprisingly, the main view of *SongSlope* was also more popular for creating playlists: Almost half of the participants (48%) used it instead of the session view (32%) or a manual approach (20%).
The application was easy to learn and use as the respective questions were both answered with a mean of 2 (1 being strongly agree) on the 7-point Likert scale. Asked about their satisfaction with *SongSlope*, participants also answered with a mean of 2.
The results of the free form fields gave us insights into the ways people used the tool. Some people confirmed the listening factors that we had also identified in the other studies: *"[I found] That I have rather short listening sessions", "I'm more the listen-to-the-whole-album person", "I listen on shuffle a lot"*[45]. They also mentioned the rediscovery use cases that we had assumed (*"I found a fiew [sic!] songs I was looking for some time ago"* [ibid.]) and one participant even confirmed our notion of the song-to-song relationship: *"Sometimes the [main] view showed me songs, which were already in my head, while listening to the actual song"* [ibid].
The results of the questionnaire therefore confirmed some of our suspicions and the high acceptance of the tool also showed that people benefited from it. But beyond these conscious self-reports we were curious whether the automatic logging in the application could teach us more about the ways people used *SongSlope*.

Results of the interaction logging

We logged four types of interaction: playlist generation, entering the session view, clicking on a song in the main view, and changing the track. Privacy-awareness nevertheless

caused a large number of listeners refraining from making their interaction logs available. In total, we received data for 249 listening sessions. We first filtered entries that contained only a single session of a single listener: We assumed that these people had only tried the add-on once before discarding it. Their behavior was characterized by exploring the application, only shortly listening to songs and trying all features before returning to their regular use of *Songbird*. Due to server problems, we additionally had to discard some of the entries and arrived at a total number of 70 sessions by 16 regular listeners in the end.
In these 70 sessions, we recorded 588 track changes (avg. 8.4 per session), 17 clicks on songs, 7 generated playlists, and 20 times where the session view was entered. One interesting thing we noted was the (relatively) low number of generated playlists and clicks on songs. We had expected to find more of that in the data, as the results from the questionnaire implied that reminiscing and rediscovery took place with the application. Looking at the low-level interaction data, however, showed the actual frequency of such events. While they happened distributed among all listening sessions, they were rather rare. Our explanation for this finding is that music listening with a media player often happens in the background: The listener selects a playlist or sets the application to shuffle, presses play and goes on to do something else. *SongSlope*'s reminiscing triggers are purely visually and therefore need the listener's attention to be useful. This explains the low number of such events in the interaction data.

6.3.4 Discussion

SongSlope gave us the chance to learn about reminiscing and rediscovery of songs in a regular listening scenario. Instead of creating an explicit application for analysis and working with a listening history, the tool was developed as an add-on for existing and popular software and was more convenient to discover and install. The large number of downloads showed that this approach had been successful.
The evaluation itself also confirmed our hypotheses regarding the usefulness of the application: *SongSlope* could help in rediscovering old favorites, explore listening sessions, and was easy to use. Especially some of the free-form answers showed the potential of the approach. Looking into the interaction on a very low level using automatic logging also gave us interesting insights into the click-to-click behavior within the application.
Overall, we were happy with the results but also found that for one, the rediscovery of songs is more of a niche topic as it is very personal and requires one's own listening history (*LastHistory* for example could also be used with an external listening history). Also, as people mainly rely on a media player for background music listening, they often simply miss what the application presents to them.
A future version could approach this problem by displaying the information not as part of a window but in a more accessible space for background information such as a screensaver. Another alternative might rely on a mobile device (where much music listening happens anyway, see section 6.1): The *Android* operating system, for example, allows

creating so-called active backgrounds, that can adjust dynamically. One option would be to show predecessors and followers there as a backdrop for any current interaction with the phone.

6.4 Summary

In this chapter, I presented three prototypes that looked in more detail into the *item* dimension of the design space. Connections between songs can arise either because of some inherent similarity (and then be consistent across all listeners) or due to some event that triggered this connection (which is only applicable for a single history). As all music listening is sequential, using these relationships can help in enriching one's interaction with music, either by helping in building playlists or rediscovering old favorites. Apart from this emphasis on the *item* dimension, the prototypes in this chapter also contained a more forward-looking approach to working with listening histories. While *Strings* or *LastHistory* provide convenient ways to explore one's past and wallow in memories, especially *Rush* shows that listening histories can also be helpful in supporting future music listening. Building playlists according to one's taste in a semi-automatic way needs not only explicit input from interaction but can also benefit from knowing more about the personal relationships between a listener and their songs. Using a former playlist in *SongSlope* for directly listening to it also lies somewhere on the border between reminiscing and playlist building.

The prospective use of music listening histories is a promising approach to helping people in a more direct way by relying on the (unobtrusive) collecting of their behavior. Recommender systems often approach the problem in an overarching, one-size-fits-all way that either uses general similarity measures for forming suggestions or does not provide enough ways to influence the results. Using a listening history for either deriving personal similarity metrics, suggesting popular playlists of one's past or helping with rediscovery in the background all support interaction with music by relying on existing data.

Apart from the benefits for prospective music use, the projects of this chapter also showed different ways to deploy these applications. A dedicated analysis tool is not necessary when a mobile application or a plug-in for a media player is much more appropriate for the task. With music listening shifting more and more towards mobile use especially deployment on handheld devices is looking to be more important in the future.

Finally, the two-fold evaluation for *SongSlope* gave more insight into the differences between perception of and actual interaction with an application. While the results from the questionnaire showed that people thought of the tool as enabling reminiscing and rediscovery, on the level of direct interaction with it the results were less clear. For seeing the use of an application a single positive experience is often enough, even when the regular use does not show the same. But for understanding all aspects of an application

having access to both sides can be very helpful.
One question that remains and has only partially been tackled with the above prototypes is how to include discovery with such personalized systems. While rediscovering old favorites is nice, a system than only allows repeating the same songs over and over again will sooner or later frustrate the listener. Therefore, integrating recommendations for discovery would be important for creating more variety while also being able to explain the suggestions with the existing listening events.

Chapter 7

Summary and Future Work

Don't wanna waste no more time.
Time's what we don't have.

– Biffy Clyro - *Living is a problem because everyone dies* –

In this dissertation, I presented an overview of music listening histories. This special data type provides deep insight into a person's taste in and behavior surrounding music. Due to music's shift towards digital means of consumption the automatic collection of such histories has become feasible and promising in capturing a large part, if not all of a person's tracks.

At the moment, the possibilities opened by music listening histories only gradually enter the consciousness of researchers and practitioners and especially that of the creators of this data. While they welcome improved music recommendations for discovery and personalized web radios, those creators are barely aware of the existence of music listening histories as all access to them only happens implicitly. Only gradually is the idea of lifelogging and the role of listening histories in it entering the public mind.

Music listening histories can be used similar to photo albums as instruments of reminiscing that allow reflecting about one's life and reliving one's past. Given the right memory triggers, the originally innocuous timestamps and song titles can turn back into parts of personal lifestories. They can also become the source for a more analytical approach that let people understand their behavior in an abstract sense, based on amounts, frequencies, and repetitions. Beyond looking backwards, however, listening histories can also act as bricks for forward-looking activities. They can form parts of future playlists and help in understanding the intrinsic and event-based similarities between two items. While listening to music, they can also form unintrusive reminders of previous favorites for rediscovery or as inspiration. The relational nature of item similarity can be made more tangible by connecting it to one's own history.

The different chapters of this dissertation shone a light on different aspects of this data type. In chapter 2, I presented an overall survey of the relations between people and music. Extensive research from musicology, human-computer interaction as well as music information retrieval has worked at unscrambling the mystery of human attraction to structured sounds. On the one hand, musicology has shown reasons for listening to music and its function in people's lives. On the other, HCI and MIR have looked at the ways in which people actually interacted with music and managed their ever-growing music collections. The shift towards digital music has led to the peculiar situation of making computer scientists experts on music consumption. I combined the research from these different fields into an overview of listening factors that influence music choice.

Chapter 3 discussed the nature of ideal and real-world music listening histories. I presented their characteristics in the context of lifelogging that aims at perfectly capturing the complete human experience and showed how listening histories are a small but relevant (and especially already available) part of a perfect lifelog. Real-world listening histories are not perfect: They lack contextual information that would be helpful for understanding listening choices and suffer from noise and gaps, representing an imperfect image of their ideal counterparts. Still, they can capture the most relevant aspects and be supplemented with data from other online services. The most suitable extension, however, proved to be their creators' minds that are able to form a coherent whole out of imperfect song lists. In chapter 3 I also presented the results of a large-scale analysis of long term histories from last.fm. This principal components analysis made use of the various variables for music listening that we identified from literature and showed their relationships and their importance. Especially seasons and the drive to music discovery proved to be distinguishing factors for listeners.

The next chapter presented visualization as a promising way to access music listening histories. As this type of data can become very complex very fast, relying on visualization as a method for aggregating abstract data into digestible pictures helps in understanding events and patterns. Various visualizations, mostly from non-researchers, are available online and present listening histories either in summarized forms that emphasize long-term trends or are good for answering single questions about the data. I presented a design space for music listening history data that consisted of the three main dimensions *time*, *items*, and *listeners*. These dimensions show the dual nature of listening either as lifelogging (*time* dimension) or relational data (*items* dimension) and also contain the possibilities opened by adding more than one history to the analysis (cf. social uses of music 2.1.3). Additional dimensions that are relevant for the resulting application contain the listener's personality and background knowledge, the target device and the purpose of the tool.

Based on these dimensions, the next two chapters give examples. The visualizations of chapter 5 focus on the *time* dimension and thus lifelogging-related tasks such as reminiscing and analysis. *Strings* and *LastHistory* explore the capabilities of timelines

for providing single listeners with access to their past listening activity. The latter extended the available data with photos and calendar events as suitable memory triggers for easier understanding. An online questionnaire showed the usefulness of the approach and provided insight into what aspect of the listening factors people were able to find using the tool. A second set of visualizations, *LoomFM* and *LastLoop* included the *listeners* dimension, either for comparing the taste of two people or for an arbitrary number of histories. Their focus was again on temporal patterns. The subsequent evaluation of *LastLoop* not only informed us about the suitability of the tool for its actual task but also helped us in learning about the impact of interface complexity on acceptance by our target audience.
Chapter 6 held another set of visualizations that explored the use of listening histories outside of a "classical" information visualization approach with a desktop-based and analysis-oriented nature. I still relied on visualization techniques for making the data understandable. *Rush* is an interaction technique for quickly creating suitable playlists using an automatic approach enhanced with listener choice. Listening histories allow deriving either overall or personalized similarity metrics. Parts of them can also be interspersed into the suggestions. The technique is also focused on mobile devices as nowadays' standard platform for music listening. *Tangle* was conceptually closest to the visualizations from the previous chapter and provided a look into the relational nature of songs in a listening history and their complex interplay. Finally, *SongSlope* is a plug-in that allows rediscovery of old favorites or using sections of a history as new playlists in the natural habitat of the listener, a desktop media player. Integration with *Songbird*'s plug-in catalogue helped us with being discovered. The combined evaluation approach of a questionnaire with low-level interaction logging uncovered the differences in perception and actual use.

In this last chapter I will discuss the contributions of this thesis (cf. section 1.4) in a more concrete way and in light of the presented prototypes. I will also show how the work presented in this thesis can be continued and how its results can help in improving the lifelogging movement and understanding people's interaction with music.

7.1 Contributions

As mentioned in chapter 1, the contributions of this thesis are relevant to researchers in information visualization and human-computer interaction (regarding the design space, visualization prototypes and listening factors), but also musicologists (regarding the data-driven methodology).

7.1.1 Deepening the understanding of human behavior surrounding music

A first contribution concerns the reasons why and the ways in which people listen to music. Musical behavior has always been a popular research topic, so several disciplines describe aspects of it. The first part this thesis contributes in this regard is an overview on different levels of detail on the relationship between people and music. The abstract *why* is answered by musicology and with the three identified reasons of *mood management* (music for regulating emotions), *cognitive functions* (music for supporting tasks and flow), and *identity construction* (music for maintaining an outward and inward image of one's personality). On a more concrete level, HCI and MIR answer *how* people listen to music. I separated music interaction into six discrete steps: *Discovery* of new music, *Acquisition* of a song, *Organization* of a music collection, *Understanding* the relationships within that collection, *Sharing* songs with other people, and *Listening* to single songs or playlists. For music listening histories especially the last step is important. While there are several studies from different disciplines that explore music listening, they differ in crucial aspects that influence the applicability of the results. I identified *realism, music granularity, context, listener information*, and *sample size* as dimensions along which the studies differ. Based on these different abstractions, I collected listening factors (see section 2.4) that describe what circumstances influence music choices in which way. While these listening factors do not represent a complete model for human music listening, they still include all factors that are relevant in music choice. Interaction designers can use these factors in deciding which aspect to emphasize and which tasks to support.

Currently, granularity in music capture mostly correlates with a smaller sample size: Music listening studies usually involve manual tasks such as reminding the participants or them noting what they are hearing at a certain point in time. Therefore, a second part of my contribution in this regard is providing a different methodology for understanding human music behavior. Music listening histories cannot only be used to help their creators or identify commonalities between them and their acquaintances but also on a larger scale. Our principal components analysis of listening histories (see section 3.4) showed a different methodology for looking at music listening. While musicological studies emphasize the manual data collection and the included capturing of valuable listening context (e.g., [120]), automatic listening history collection as provided by various online services has the advantage of providing large numbers of histories on the very detailed level of single songs. Both work as supplements to each other, but so far musicology is only reluctantly approaching large data studies (cf., [72]). Using a large sample of automatically collected listening histories we were able to categorize 48 different variables identified in music literature. We were able to sort them by their respective importance in musical choice and collect related variables into components. This not only makes it easier to find connections between musical behavior but also shows relations between different behaviors that might not have been obvious before. Such approaches are promising not only for future musicological studies but also future humanitarian studies in general. The ever-growing compound of data from social network- and online lifelogging services enables understanding human behavior on a more detailed scale than ever before.

7.1.2 Structuring the design space and suggesting ways for visualizing personal temporal data

My second contribution concerns music listening histories as data sets. While the previous contribution was about the analysis of human music interaction in general, chapter 3 discussed the characteristics of listening histories. For one, I presented listening histories as part of lifelogging data and showed the similarities regarding collection, access and use. Based on the listening factors identified in chapter 2, I created a mapping of how these factors should be represented in an ideal listening history. I also gave an overview of the state and problems of real world listening histories and different services that provide such data.

The design space for this data type (see chapter 4) based on the three main dimensions *time*, *items*, and *listeners* is not only useful for structuring listening histories but is abstract enough to be applied to various types of lifelogging data (see below, section 7.2). The additional dimensions (*background, device, task*) are equally flexible and can be applied to different use cases with related data. Despite the abstraction, the design space nevertheless allows having a clearer understanding of a use case and identifying shortcomings of current systems (see for example figure 4.4).

The dissertation also contains various concrete examples how music listening histories can be visualized or integrated into suitable use cases. Chapters 5 and 6 show different approaches that demonstrate all of the main dimensions and also contain examples for different devices and tasks. Apart from the design and the motivations behind these prototypes, the different evaluation methods that we chose are also valuable for future research. Online questionnaires, lab studies and low-level interaction logging unearth different aspects of the experience and can be combined to arrive at more complete pictures.

The prototypes were developed for the only vaguely defined group of last.fm profile owners, which meant that we could not rely on any background knowledge besides trivialities like song titles and dates. Gradual discovering of functionality, adding suitable contextual data and using minimal and thus non-threatening interfaces, however, let us support different types of background knowledge, from available personal memories, to knowledge about music or visualization techniques. Especially the concept of immediate benefits of the static visualization was enticing to people who discovered our tools online and caused them to dig deeper.

7.1.3 Using visualization approaches for music-related tasks

Visualization techniques are not only viable for explicit visualization and analysis tasks. This dissertation also showed how they could be used for tasks beyond gaining insight. While the prototypes presented in chapter 5 were more about "classical" information visualization tools, the next chapter (chapter 6) showed how this type of data could be used in other tasks and how visualization techniques can be applied to non-analysis tasks. When using listening histories as indicators for the relational aspects of music (see chapter 6), visualization techniques such as trees (*Rush*) and graphs (*Tangle*) can be applied in a non-analytical context and work for reducing data complexity just as well. Non-canonical visualizations can also refer to non-desktop environments. The preliminary studies for *Rush*, for example, showed the impact of reduced screen space, inaccurate input methods such as touch-screens, and occlusion for shaping the interaction. Integrating visualization techniques into such scenarios requires suitable adaptations and recent work on visualizations on tabletops (e.g., [170]) explored such characteristics. *Rush* goes into a similar direction for visualizations on a smaller scale with touch-screens. Similarly, *SongSlope* tackled the problem of a background visualization that was not in the focus of attention of the listener. This comes with its own set of problems and challenges.
Additionally, I discussed the concept of song similarity in this thesis and showed how co-occurance can represent more than just overarching relatedness. So far, musical co-occurance is mainly used for automatically deriving inherent similarities using collaborative filtering, but can contain interesting connections for a personal history. In this thesis I showed ways to make such connections visible (clicking on a song in *Strings* to let all instances collapse into one, multi-song arcs in *LastHistory*) and explore them (e.g., unintrusively in *SongSlope*). Beyond that, especially the gradual change from connections within a single listening history to those within a whole set of them could be interesting, for example, for explaining recommendations.

7.2 Listening Histories in the context of lifelogging

While this dissertation discussed the topic of music listening histories in detail and presented suitable visualization approaches, they are not limited to this type of data. As listening histories are a section of lifelogging data they are also similar in their characteristics: this allows applying discovered aspects and methods also to other such data types. In this section I will discuss how to generalize the results of the thesis to other lifelogging data (cf. [13]).
The focus on time in lifelogging data (see section 3.1) ensures that the main element of the design space is applicable to general lifelogging data. In *LastHistory*, for example, we added photos, calendar entries and music events to the same timeline. Both lifelogging data and autobiographical memory are centered around the concept of time

which enables putting things into context, perspective and seeing connections. Therefore, time is the dimension that ties all lifelogging data together and putting the focus of a visualization on it becomes even more important when multiple data sources are integrated. Facebook's recently announced Timeline feature[1], for example, will provide such a general-purpose timeline and connect data points from various sources (automatically and manually created) to a central time bar. Apart from this most extreme example, almost all online services that contain personally created data rely on a time-centric presentation for navigation. The *time* dimension, however, should also provide some way of abstraction to uncover patterns on different levels of detail. The highest overview might show general overarching trends that span the whole lifetime (as an example from music: a streamgraph [27]). The lowest level shows every single item, their co-occurances and other connections. In between, automated summaries are created for arbitrary timeframes depicting amounts, trends and items receiving the most attention (e.g., songs that are repeatedly heard).
The second dimension of *items* is flexible enough to represent any type of (possibly repeated) items. A major distinction can be drawn between items that were actively created by the owner of the lifelog (e.g., photos, blog posts), those passively consumed (e.g., songs, books), and contextual information (e.g., places, social surroundings). Especially when relying on memories for making sense of the log, the generation effect in actively created items and the context are helpful for triggering memories. Depending on the type of item, its duration can also be much more important than in the music case. In all of the presented prototypes in this dissertation I ignored this aspect, as all songs have roughly the same length of four minutes. For other media that are more varied in length (e.g., movies) and those that even require multiple sessions (e.g., books) other ways of presentation would be more suitable. Similarly, parts of the contextual information are usually overlapping (e.g., location and mood), available in clearly defined states (e.g., asleep/awake), and work more as backdrops and therefore should be presented in different ways than the other items. How to distinguish different types of items on a common timeline also is important. Options are spatial separation (e.g., different streams in *LastHistory*), color coding (e.g., genres in *LastLoop*) or discovery through interactive exploration (e.g., hovering over a song to get the actual genre in *LastHistory*). Just as with *time*, *items* also needs suitable ways of abstraction to make larger sections of the data understandable. Usually, item types provide forms of hierarchy (e.g., genres in media items, status updates/messages/blog posts for personally created text) that can be leveraged in this regard. Automatic summaries of certain timeframes should also allow for abstraction on the *items* dimension. Another way to reduce the complexity of the data is filtering. Depending on the task only certain items or types of them might be shown. While building a playlist, context is not important (*Rush*) and for the currently playing song only surrounding listening sessions are relevant (*SongSlope*). For general lifelogs, the notion of co-occurance and subsequent relatedness extends to connections between different types of items, merging for example vacations with songs and people

[1] http://www.facebook.com/about/timeline

with places, thus resulting in a much richer network of connections. Depending on the task, such cross-type connections can be used for reminiscing (e.g., displaying photos of a trip while playing suitable music as in *LastHistory*) or for more creative recommender systems (e.g., a system that suggests music based on the current location).
The last data space dimension, *listeners*, describes ways to merge more than one person's history and would be more appropriately named *lifeloggers* in this regard. Displaying several lifelogs in full fidelity is difficult without creating overly complex visual representations. Depending on the number of histories and the task, filtering or summarizations can help in reducing this complexity. For comparisons, focusing on one history (mostly the analyst's) and putting it in context to the others should work best. It also makes a difference which part of the history is compared. Different types of lifelog data require different approaches for comparison: Regular items such as songs or other passively consumed media can be compared based on changes in quantities and taste. Actively created items such as textual data are promising targets for identifying keywords and their frequencies or more sophisticated types of text analysis (e.g., sentiment analysis). Finally, comparing contextual information such as locations or social surroundings might be easier to do using maps and graph visualizations than nodes connected to timelines. Suitable aggregation of multiple lifelogs can take different forms, depending on the differences between them (e.g., highlighting the small differences in otherwise very similar logs) and the relationships between their creators (e.g., groups of friends or strangers). Co-occurance of items can also happen in multiple logs at the same time, thus hinting at either a very similar taste (or large coincidences) or time spent together. Repeating sequences that appear throughout multiple histories show effects of external sources such as a predefined album order, or the propagation of playlists.
All in all, the presented design space can also be applied to general lifelogging data. Presented visualization concepts (in chapters 5 and 6) are also applicable to other types of data, depending on task and complexity of the data.

7.3 Future Work

A thesis can only be a first step into a certain direction. Research efforts especially in the lifelogging area are still at their beginning. The amount of collected data is ever-growing, privacy concerns and questions of ownership are still open, and what to do with lifelogging data once it has become available has also not been answered conclusively. The work presented in this dissertation emphasized two general directions: How to use available lifelogging data in research in the humanities and how to support the owners of this data in their everyday lives.

7.3.1 Understanding human behavior through lifelogs

The ever-growing amounts of lifelogging data allow understanding behavior on a much larger and at the same time more detailed level. Instead of taking small samples and manually collecting data (e.g., [120]), researchers can just rely on the data that is readily available on online platforms. Our own study of listening behavior (see section 3.4) let us categorize variables of listening behavior and estimate their importance for listening histories of multiple years, something which would not have been possible using more traditional means of data collection. For music listening alone, the available data allows analyses on very detailed levels without ever leaving the lab. The current lack of contextual information is still a problem but as growth in this regard would be positive not only for researchers but also companies (with Facebook and Google being the most eager candidates) it is only a matter of time until collection of such information is integrated with existing services. Having close to complete lifelogs available would allow researchers to actually understand human music listening and create a model that predicts music choice (similar to [178]) or even general behavior.
The 'big data' approach [4] also enables learning about data without forming hypotheses beforehand. The patterns within the data can be uncovered without knowing anything about the depicted scenario which even machines are able to do. This can uncover completely unexpected patterns within a data set. Once the data set has reached a certain size, even translating languages in a statistical way works without knowing anything about the contents.
Large lifelogging data sets will play a dominant role in medicine. While our understanding of the human body right now is primarily based on anecdotes, experiments and abstract statistics, having detailed lifelogs available will revolutionize predictions for proneness to diseases and recommendations for a healthy lifestyle. Simply logging the meals throughout a day could make identifying different types of eaters much easier and understanding the connections between food intake and health characteristics.
Finally, understanding human behavior will also happen on a personal level. Based on the identified characteristics people will be eager to compare themselves to the norm and learn about their own quirks.

7.3.2 The need for personal visualization

The amount of personal data available online will only increase in the future. Given the growing number of small sensors and personal devices and the ease of automatic collection of it will lead to an intensely data-rich future. However, it is still not clear if this future will have a place for access to one's own data. At the moment, uses of personal data are mostly implicit and intransparent. Webradios, recommender systems and playlist generators just "work" or they don't. Even if all activities are based on collected behavior data, this fact is not obvious to the person sitting in front of the computer. They might not even be aware that any data collection is taking place. Yet, this collection is

necessary for providing most of the interesting, new services that the web 2.0 brought and companies are not necessarily keen on sharing this data and enabling switching to competitors. Political regulations for making personal data available will, however, gain no tractions without a significant amount of backing in the population.
One way to create interest in personal data is giving people the right tools to work with their own data. As discussed above (see introduction to chapter 5), practically everybody is interested in their own lives. But right now, only researchers and hobbyists knowledgeable in statistics and visualization are able to draw something from this type of data. Beyond recommender systems or manual access to single items, the average creator of this type of data has no way of working with it. Therefore, practitioners, researchers and the companies collecting this data have to work on ways to give the data back. These are the blacksmiths of the data age that create the instruments for everybody to work with their own data. Personal visualizations are one type of example that enables understanding and finding patterns and in general to use lifelogging data to the fullest of the five Rs [149]. More and more such tools will be necessary in the future. Research areas in this regard are for one improving the access to these patterns, providing explanations, and keeping interfaces clear and easy to use. The focus on online data also enables new forms of collaboration between visualization creators and visualization audiences: Direct feedback becomes available the moment a new visualization goes online which drastically shortens the development cycles and has the potential to understand what people need better and give them what they want. Finally, the transfer of visualizations to handheld touch-screen devices will increase our understanding of the value of these techniques and how to keep them manageable even on the go.

7.4 Conclusion

Lives are nothing but long sequences of moments. Each moment has its place in the long chain of causality that forms our characters and life situations and we would not be where we are with a single one of them missing. The collection of these moments makes us what we are and as beings with a firm sense for the past we cannot spend our lives without looking back sometimes. Naturally, we create mementos, artifacts that help us in remembering and reliving moments from the past. Writing let us carry our ideas through time, just as photos enabled capturing one frail perspective of a situation. At the beginning of the age of lifelogging, where our machines have learnt to remember for us, we have no longer single, analog things to remind us, but rich and extensive databases of digital memories. But however detailed and precise these databases will become, without their creator as observer they will be stale and fake, like flipping through a stranger's photo album. Only a lifelog's owner will be able to turn all bits and bytes back into smell and taste and heartbeat.

Appendices

Publication	Project	Corresponding section(s)
D. Baur, and A. Butz. Pulling Strings from a Tangle: Visualizing a Personal Music Listening History. In *Proc. IUI '09*, ACM, 2009, pp. 439-444	*Strings* and *Tangle*	Section 5.1 and section 6.2
D. Baur, S. Boring, and A. Butz. Rush: Repeated Recommendations on Mobile Devices. In *Proc. IUI '10*, ACM, 2010, pp. 91-100	*Rush*	Section 6.1
Y. Chen, **D. Baur**, and A. Butz. Gaining Musical Insights: Visualizing Multiple Listening Histories. In *Workshop on Visual Interfaces to the Social and Semantic Web (VISSW2010)*, 2010	*LoomFM*	Section 5.3
D. Baur. Visualizing Media and Music Histories. In *Know Thyself: Monitoring and Reflecting on Facets of One's Life (in conjunction with CHI 2010)*, 2010	*Strings*, *Tangle*, *LoomFM*	see above
D. Baur, F. Seiffert, M. Sedlmair, and S. Boring. The Streams of Our Lives: Visualizing Listening Histories in Context. In *IEEE Transactions on Visualization and Computer Graphics, vol. 16, no. 6*, IEEE, 2010, pp. 1119 -1128	*LastHistory*	Section 5.2
D. Baur, J. Büttgen, and A. Butz. Listening Factors: A Large-Scale Principal Components Analysis of Long-Term Music Listening Histories. To appear in: *Proc. CHI 2012*, ACM, 2012	*Last.fm PCA*	Section 3.4

Table A.1 List of peer-reviewed publications with results from this thesis.

A Publications

This dissertation is based on several papers that have been published in peer-reviewed conferences and workshops (see table A.1).

B Dissertation listening history

Due to the topic of this dissertation, I felt required to keep track of the music that I listened to while writing it. During the four months (June to September 2011) I have listened to a total of about 1,300 songs. The following figures show different aspects of this data, using some of the visualizations presented here. Last.fm automatically calculates some statistics from the data (favorite tracks see figure B.1, favorite artists see figure B.2). Visualizations from chapter 5 mainly show temporal aspects: *Strings* displays listening sessions (see figures B.4 and B.5), *LastHistory* and *LastLoop* show daily rhythms (see figures B.6 and B.3), and LoomFM uncovers repeating patterns (see figures

B.9). *Tangle*, which has been presented in chapter 6, shows song to song connections in a graph (see figures B.7 and B.8).

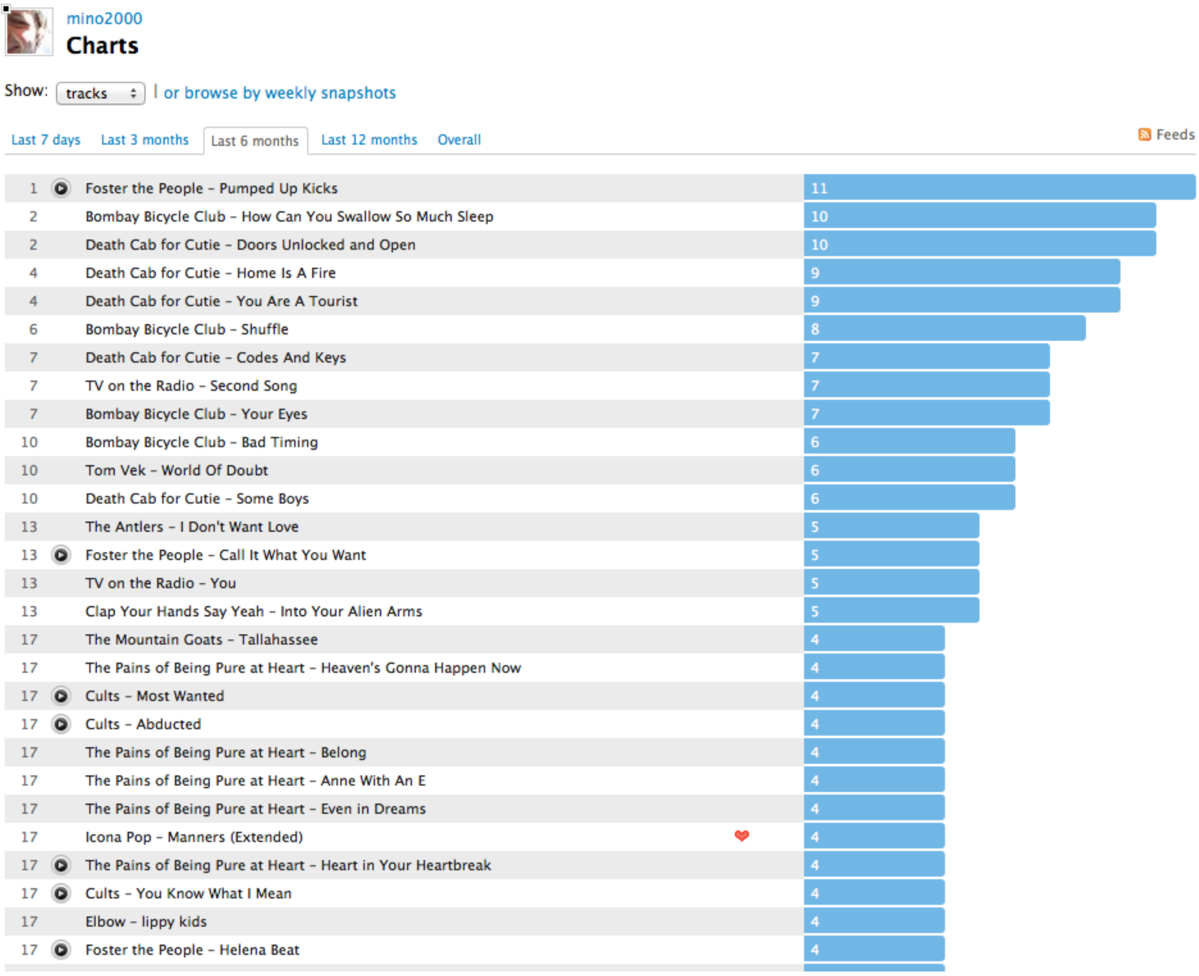

Figure B.1 Songs listened to while writing, sorted by favorite tracks (source: last.fm).

C Rotated component matrix from the large-scale last.fm PCA

In section 3.4 I presented the results from a principal components analysis of last.fm listening histories. After deriving general statistics from roughly a 1,000 listening histories, we took a sample containing 310 long-term (i.e., more than two years duration) histories. Based on them we determined thirteen principal components from 48 variables. Figure C.10 shows the resulting rotated component matrix and the connections between factors and variables.

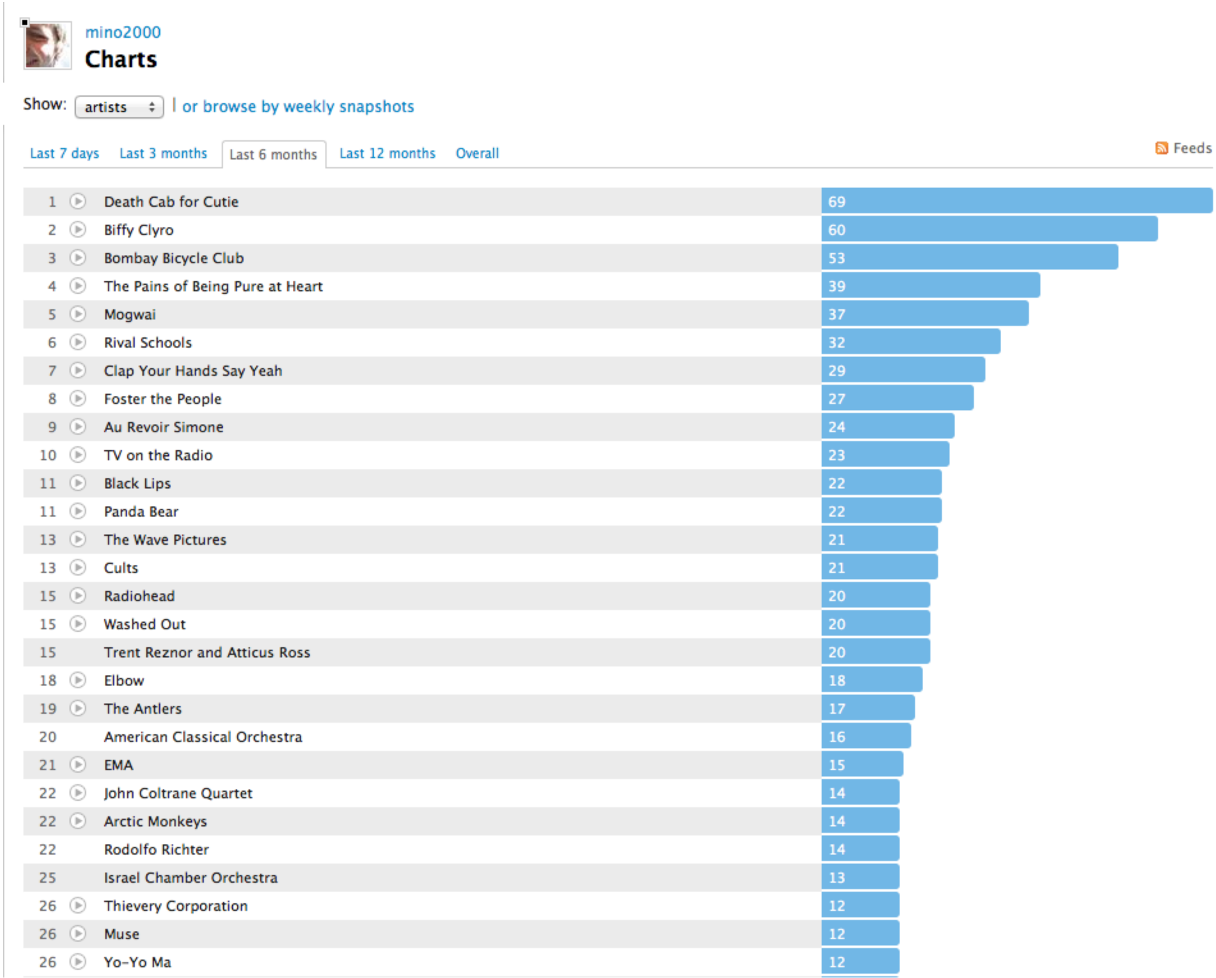

Figure B.2 Songs listened to while writing, sorted by favorite artists (source: last.fm).

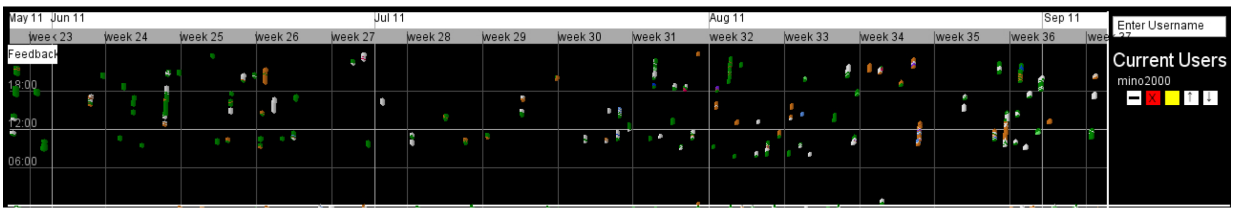

Figure B.3 *LastLoop* uncovers sleep wake rhythms, similar to *LastHistory*.

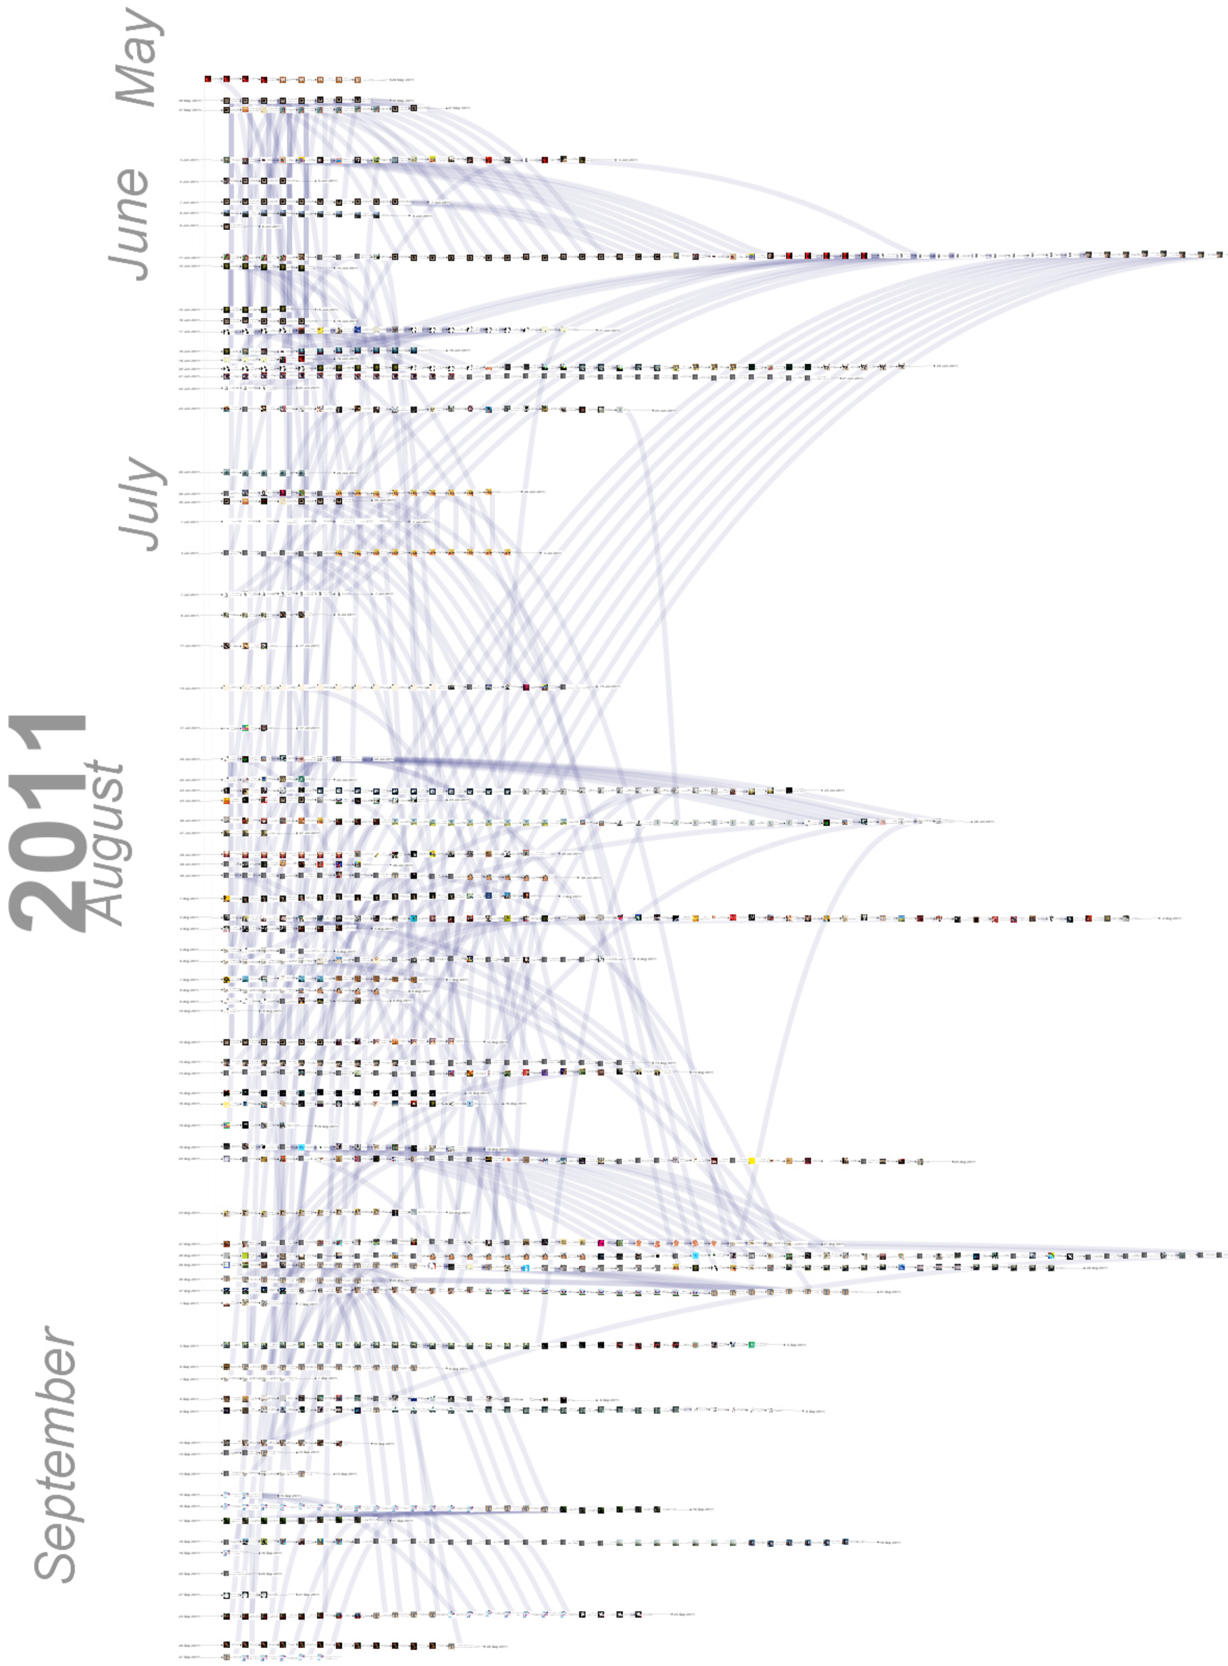

Figure B.4 The same songs, organized by listening sessions in *Strings*.

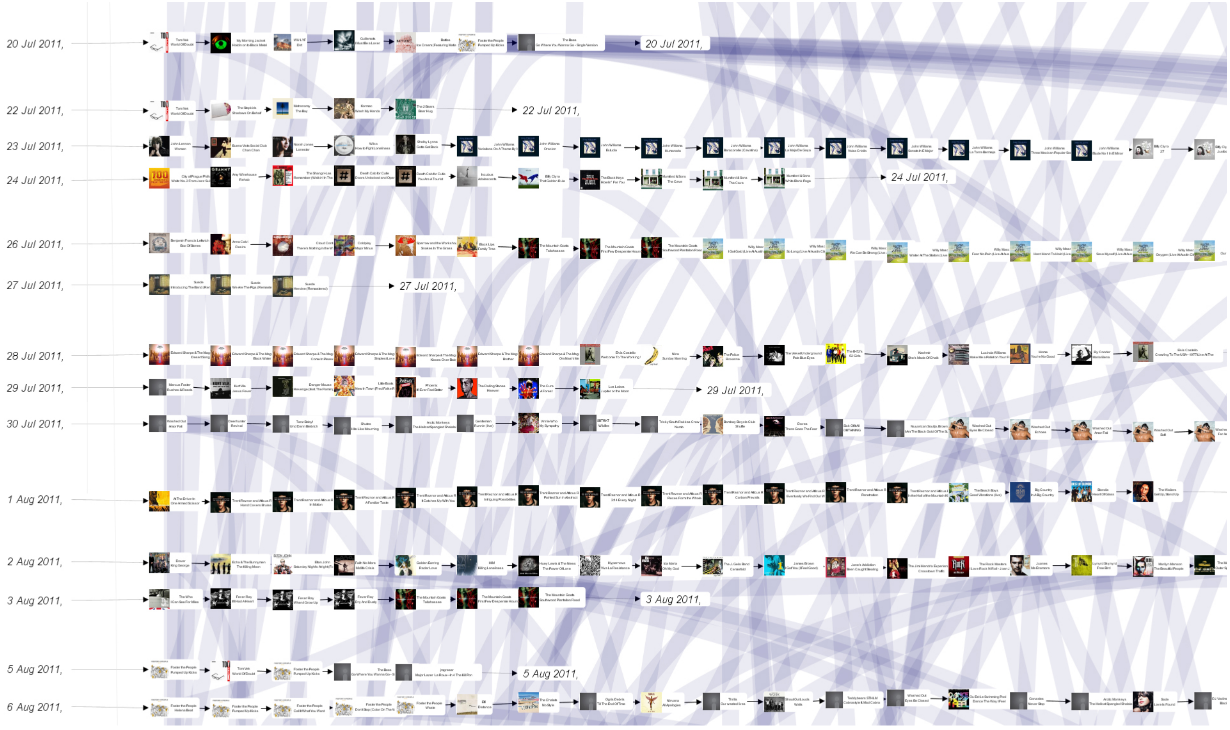

Figure B.5 Detail from figure B.4.

Figure B.6 The two-dimensional timeline in *LastHistory* shows sleep and wake rhythms.

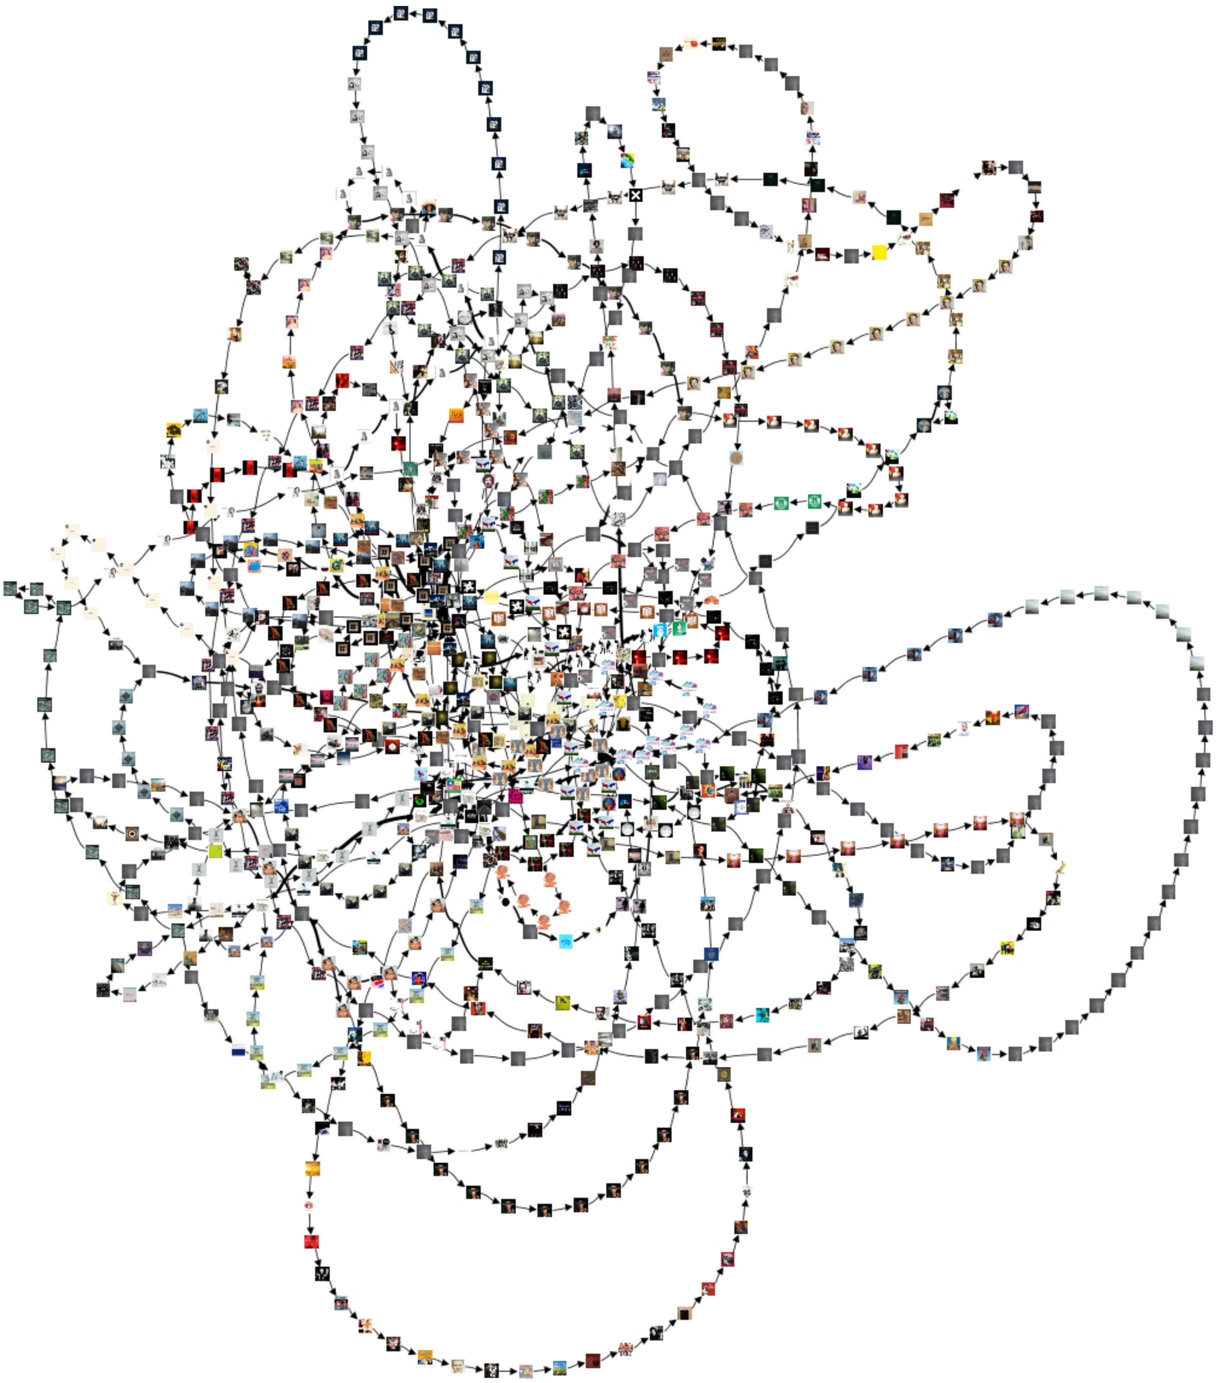

Figure B.7 *Tangle* shows sequential connections between tracks.

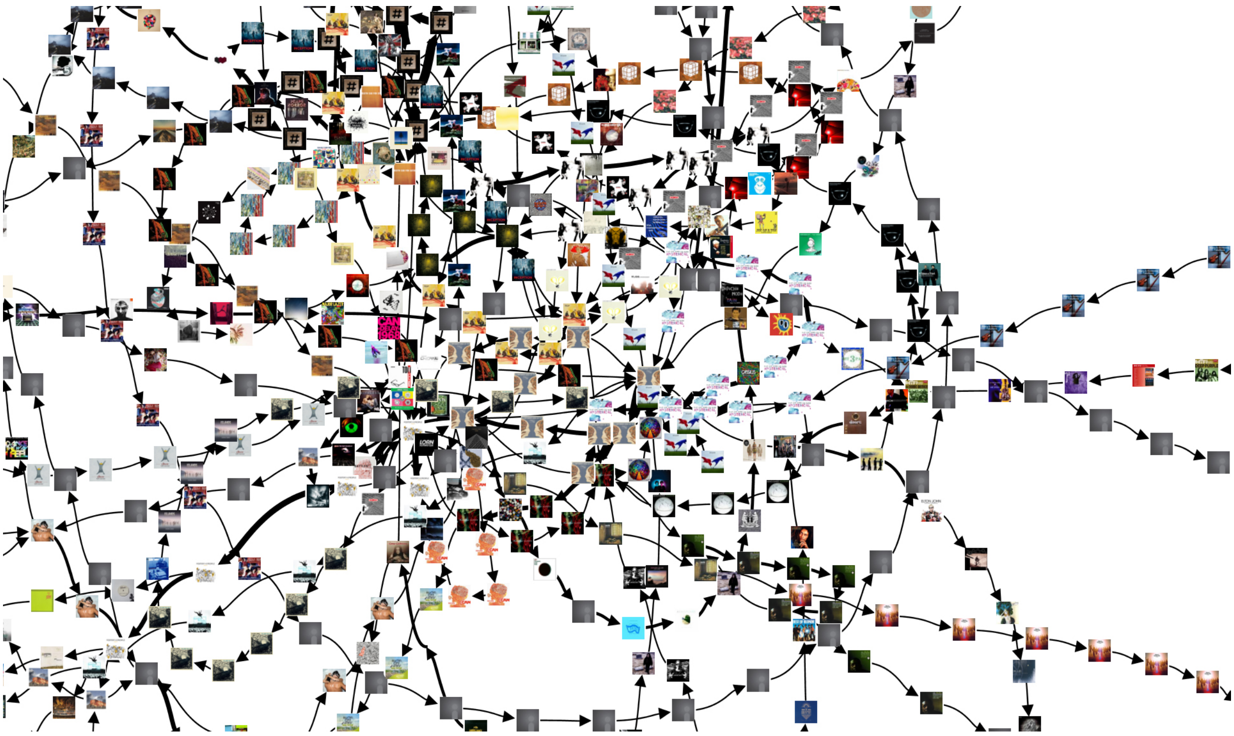

Figure B.8 Detail from figure B.7.

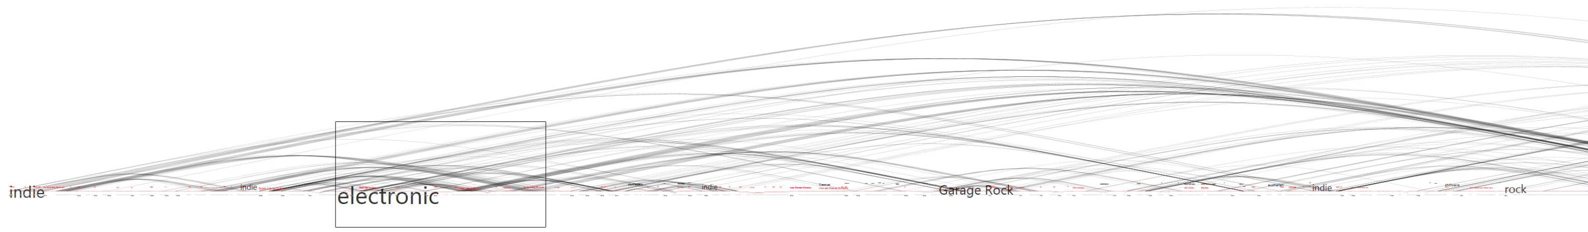
indie
indie
electronic
indie
Garage Rock
indie
rock

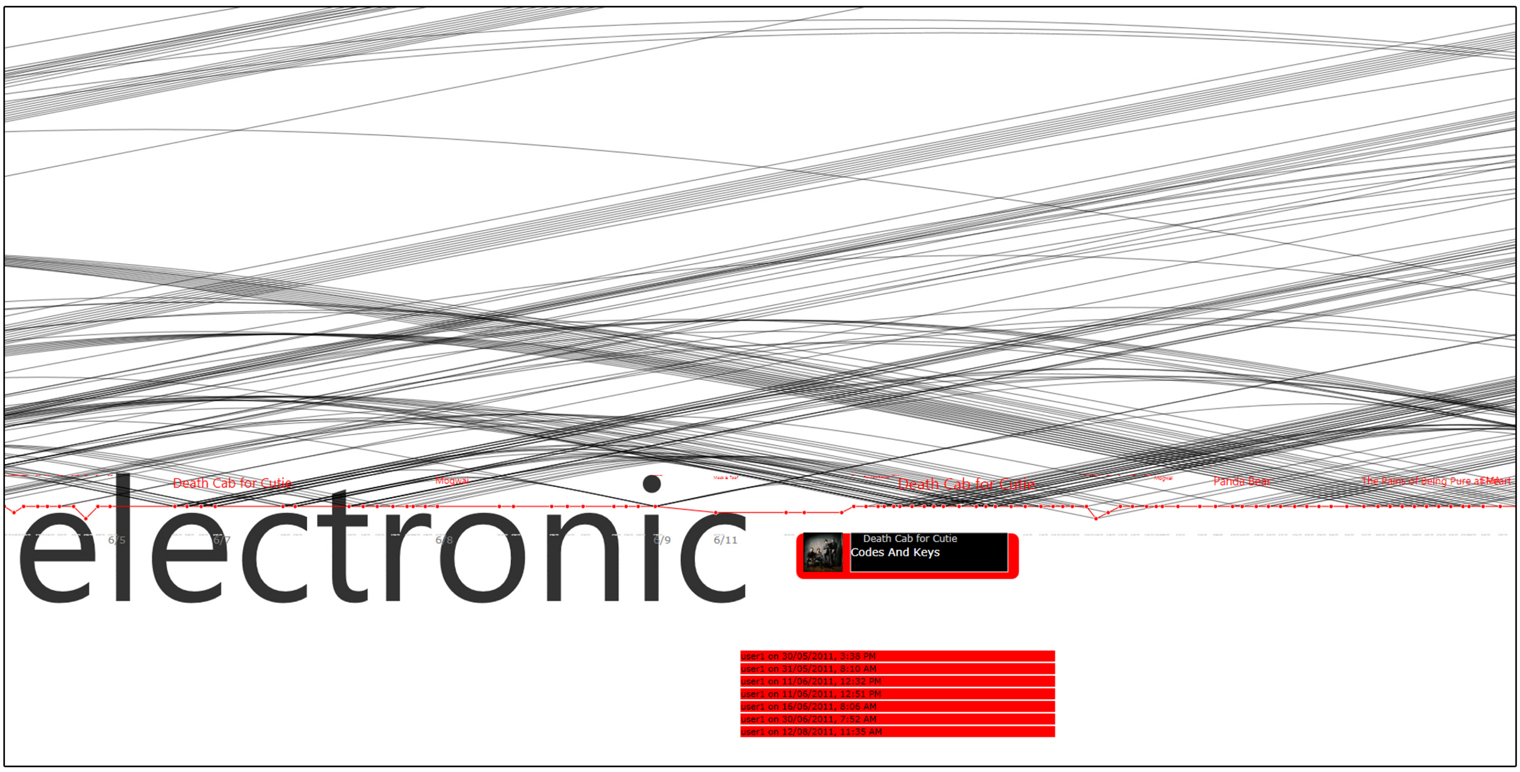
Death Cab for Cutie
Mogwai
Death Cab for Cutie
Panda Bear
electronic
6/5
6/7
6/8
6/9
6/11
Death Cab for Cutie
Codes And Keys
user1 on 30/05/2011, 3:38 PM
user1 on 31/05/2011, 8:10 AM
user1 on 11/06/2011, 12:32 PM
user1 on 11/06/2011, 12:51 PM
user1 on 16/06/2011, 8:06 AM
user1 on 30/06/2011, 7:52 AM
user1 on 12/08/2011, 11:35 AM

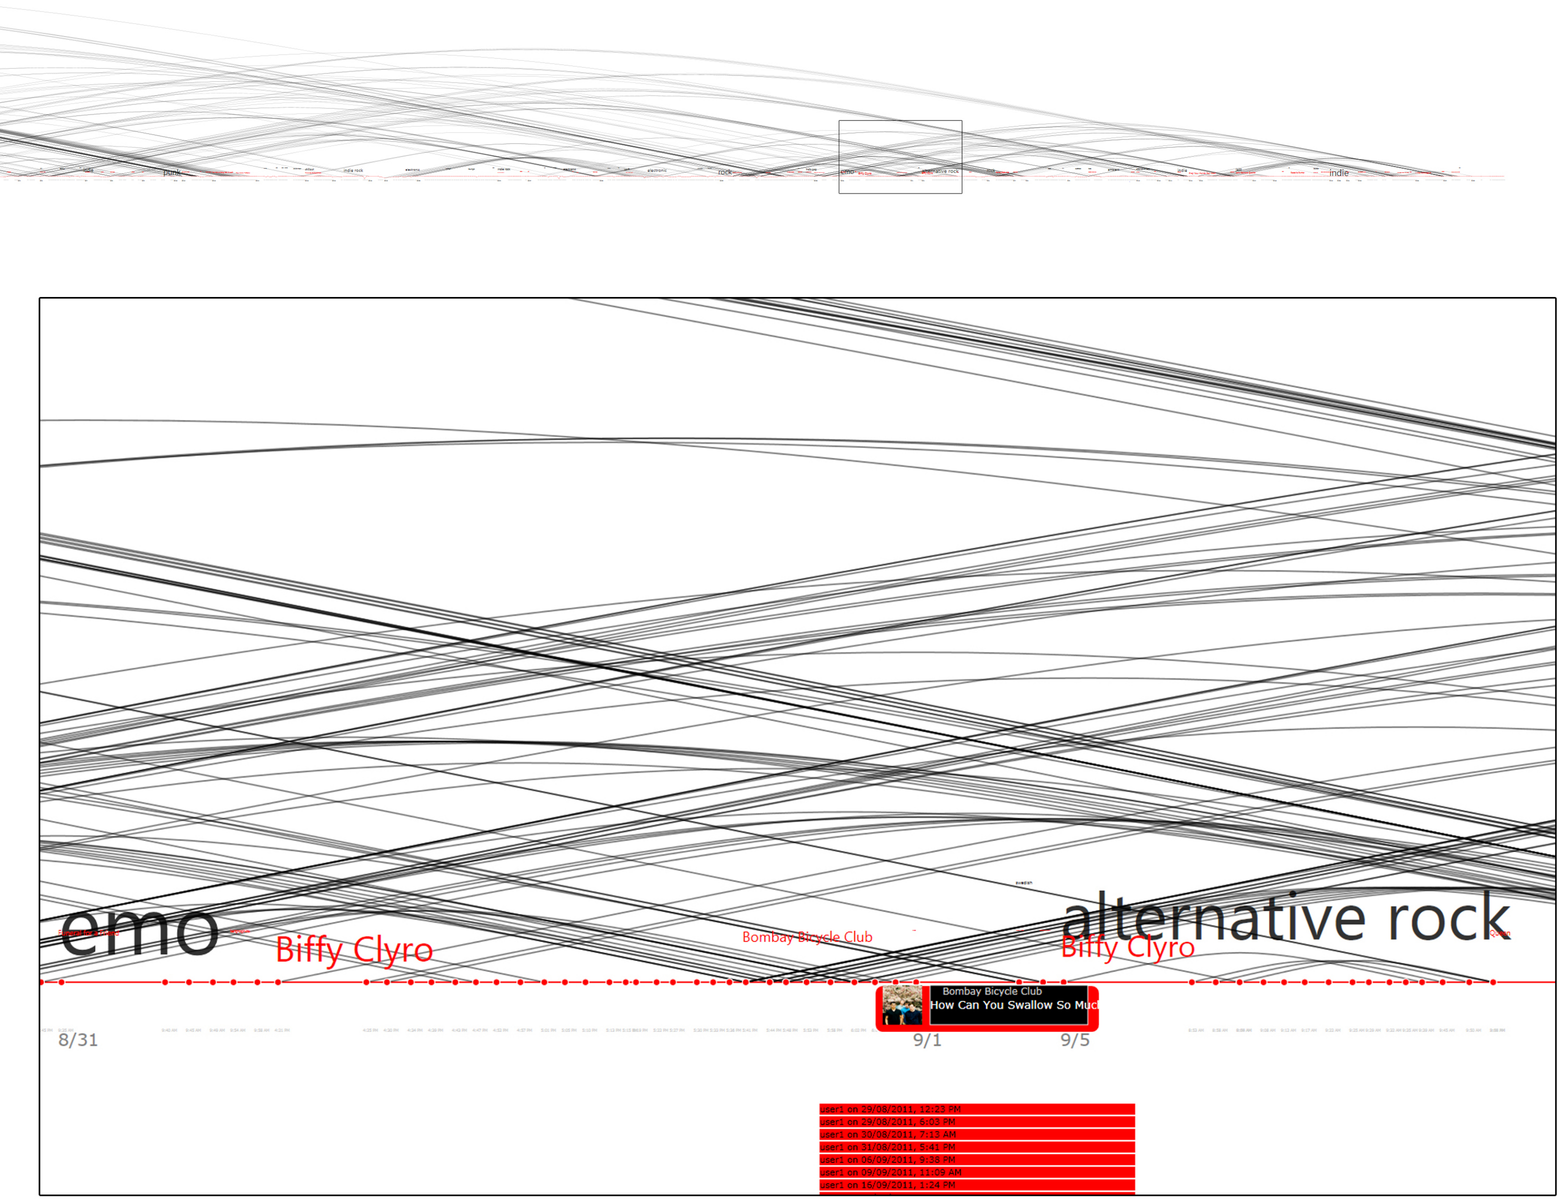

Figure B.9 *LoomFM* uncovers the general repetitive structure of the dissertation songs.

Rotated Component Matrix[a]

	Component												
	1	2	3	4	5	6	7	8	9	10	11	12	13
artUni_hlfyr2	,902						,154						
artUni_aut	,835												
gnrUni_hlfyr2	,831	-,199		-,162					,251				
gnrUni_aut	,750	-,213		-,279					,234				
artUni_sum	,667		,284								,224		
artUni_night	,595		,475	,257		,189	-,150			-,319	-,209		-,193
gnrUni_sum	,571	-,229	,226						,146		,293		
distinct_art	,566	-,526	,189	-,156	-,247		,267			-,168			
gnrUni_night	,544	-,206	,457				-,159		,259	-,364	-,153		-,194
Onetimer	,540	-,535	,259	-,206	-,193	-,160	,308	,163					
artUni_day	,524	-,180	,208	-,243			,522			,207	,221		,192
perma_art		,831		,309				,148					
perma_alb		,822		,251					-,218				
compl_alb		,751				-,163	,151	,201					
distinct_alb	,460	-,609	,193	-,257	-,224		,254			-,192			
distinct_sng	,516	-,550	,242	-,243	-,289		,286	,171		-,151			
perma_s		,537			-,206			-,356		-,220	,270		
balanceY	,300	-,393		-,342		-,315				-,252			,169
artUni_hlfyr1	,367		,783				,338						
gnrUni_hlfyr1	,223	-,226	,774				,293		,228				
y_winter	-,304		,715	-,168					-,182				
genreUni_spr	,315	-,282	,713						,218				
artUni_spr	,444		,696					-,151					
genderAvail		,208		,824									
countryAvail		,179		,814									
ageAvail	-,185	,179		,763									
sessionLg_s				-,221	,783	-,216							
introsong	-,161	,285		,259	-,611	,280	-,176			,152			
tracksDay	-,314	,372	-,193		,560	-,348						,218	
playcount	-,276	,299	-,170	,177	,530	-,391				,267		,169	
shuffle_art				,192	,529	,236		,282					,223
loyal_genre						,914							
loyal_art						,879							
artUni_win	,384		,311				,685						-,190
gnrUni_win	,280	-,152	,322				,602		,305	-,185			-,200
gnrUni_day	,416	-,409	,163	-,348			,418		,246	,215	,256		
shuffle_alb								-,816					
discover_art	-,381			-,148				,522	,181	,257			
subscriber								-,147	,731		-,160		
Online	,316		,213	-,213	,165			,183	,537				
Active				,244		-,199	-,158			,635	-,276		
songLg	,281				-,197	,165				,464			-,252
Playlists	,162			,155	,155						,704		
inorder_alb		,227		,216	,261			,405	-,149	,167	-,426		
Friends												,820	
Shouts		,154			,203							,742	
Skips							-,180						,813
inorder_art		,184	-,238	,172		,170	,336				-,210		,513

Figure C.10 Rotated component matrix from the large-scale last.fm analysis showing thirteen components and their connections to the underlying variables (source: [26]).

BIBLIOGRAPHY

[1] ADJEI, C., AND HOLLAND-CUNZ, N. *Monitoring and Visualizing Last.fm*. Diplom thesis, Fachhochschule Mainz, 2008.

[2] ADOMAVICIUS, G., AND TUZHILIN, A. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering 17*, 6 (June 2005), 734–749.

[3] AMAR, R., AND STASKO, J. A Knowledge Task-Based Framework for Design and Evaluation of Information Visualizations. In *IEEE Symposium on Information Visualization* (2004), pp. 143–149.

[4] ANDERSON, C. The End of Theory: The Data Deluge Makes the Scientific Method Obsolete. *Wired 16*, 07 (2007).

[5] ANDRIC, A., AND HAUS, G. Estimating Quality of Playlists by Sight. *First International Conference on Automated Production of Cross Media Content for Multi-Channel Distribution (AXMEDIS'05)* (2005), 68–74.

[6] APTER, M. J. *Reversal theory: Motivation, emotion and personality*. Taylor & Francis, Florence, KY, US, 1989.

[7] AUCOUTURIER, J.-J., AND PACHET, F. Representing Musical Genre: A State of the Art. *Journal of New Music Research 32*, 1 (2003), 83–93.

[8] BAINBRIDGE, D., CUNNINGHAM, S., AND DOWNIE, J. How people describe their music information needs: A grounded theory analysis of music queries. In *Proc. ISMIR* (2003), no. Table 2, pp. 221–222.

[9] BARABASI, A. The origin of bursts and heavy tails in human dynamics. *Nature 435* (2005), 207–211.

[10] BAUR, D., BORING, S., AND BUTZ, A. Rush: Repeated Recommendations on Mobile Devices. In *Proc. IUI '10* (2010), ACM, pp. 91–100.

[11] BAUR, D., BÜTTGEN, J., AND BUTZ, A. Listening Factors: A Large-Scale Principal Components Analysis of Long-Term Music Listening Histories. In *CHI 2012* (2012).

[12] BAUR, D., AND BUTZ, A. Pulling strings from a tangle: visualizing a personal music listening history. In *Proc. IUI '09* (2009), pp. 439–444.

[13] BAUR, D., SEIFFERT, F., SEDLMAIR, M., AND BORING, S. The Streams of Our Lives: Visualizing Listening Histories in Context. *IEEE Transactions on Visualization and Computer Graphics 16*, 6 (2010), 1119–1128.

[14] BEDERSON, B. B., AND HOLLAN, J. D. Pad ++ : A Zooming Graphical Interface for Exploring Alternate Interface Physics. In *Proc. UIST '94* (1994), pp. 17–26.

[15] BELL, G., AND GEMMELL, J. *Total Recall: How the E-Memory Revolution Will Change Everything*. Dutton, New York, New York, USA, 2009.

[16] BENTLEY, F., METCALF, C., AND HARBOE, G. Personal vs. commercial content: the similarities between consumer use of photos and music. In *Proc. SIGCHI conference on Human factors in computing systems* (2006), pp. 667–676.

[17] BHATTACHARJEE, S., GOPAL, R. D., LERTWACHARA, K., MARSDEN, J. R., AND TELANG, R. The effect of digital sharing technologies on music markets: A survival analysis of albums on ranking charts. *Management Science 53*, 9 (2007), 1359–1374.

[18] BISCHOFF, K., FIRAN, C. S., NEJDL, W., AND PAIU, R. Can all tags be used for search? In *Proceeding of the 17th ACM conference on Information and knowledge management* (2008), ACM, pp. 203–212.

[19] BLUM, M., PENTLAND, A., AND TROSTER, G. Insense: Interest-based life logging. *Multimedia, IEEE 13*, 4 (2006), 40–48.

[20] BOAL-PALHEIROS, G. M., AND HARGREAVES, D. J. Listening to music at home and at school. *British Journal of Music Education 18*, 02 (July 2001), 103–118.

[21] BPI. Digital Music Nation 2010. Tech. rep., BPI - The British Recorded Music Industry, London, UK, 2010.

[22] BROWN, B., GEELHOED, E., AND SELLEN, A. The use of conventional and new music media: Implications for future technologies. In *Human-computer interaction: INTERACT'01: IFIP TC. 13 International Conference on Human-Comupter Interaction, 9th-13th July 2001, Tokyo, Japan* (2001), IOS Press, pp. 67–75.

[23] BROWN, B., SELLEN, A., AND GEELHOED, E. Music sharing as a computer supported collaborative application. In *ECSCW 2001* (2001), Springer, pp. 179–198.

[24] BRYANT, W. *Virtual music communities: The folk-music internet discussion group as a cultural system*. Ph.d., University of California, Los Angeles, 1995.

[25] BUSH, V. As we may think. *The Atlantic 3*, 2.

[26] BÜTTGEN, J. *What's in a history? A large-scale statistical analysis of Last.FM data*. Diplom thesis, Ludwig-Maximilians-Universität, München, 2010.

[27] BYRON, L., AND WATTENBERG, M. Stacked graphs-geometry & aesthetics. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (2008), 1245–1252.

[28] CARD, S., AND MACKINLAY, J. The structure of the information visualization design space. In *Proceedings INFOVIS '97 Proceedings of the 1997 IEEE Symposium on Information Visualization (InfoVis '97)* (Washington, DC, USA, 1997), IEEE Computer Society, pp. 92–100.

[29] CARD, S. K., MACKINLAY, J. D., AND SHNEIDERMAN, B. Chapter 1 - Information Visualization. In *Readings in Information Visualization. Using Vision to Think*, S. K. Card, J. D. Mackinlay, and B. Shneiderman, Eds. Morgan Kaufmann, 1999, ch. 1, pp. 1–34.

[30] CELMA, O. *Music Recommendation and Discovery in the Long Tail*. Springer, Berlin / Heidelberg, Germany, 2010.

[31] CHAMORRO-PREMUZIC, T., AND FURNHAM, A. Personality and music: can traits explain how people use music in everyday life? *British journal of psychology 98*, Pt 2 (2007), 175–185.

[32] CHAMORRO-PREMUZIC, T., SWAMI, V., FURNHAM, A., AND MAAKIP, I. The Big Five Personality Traits and Uses of Music. *Journal of Individual Differences 30*, 1 (2009), 20–27.

[33] CHEN, Y.-X., BAUR, D., AND BUTZ, A. Gaining Musical Insights : Visualizing Multiple Listening Histories. In *Workshop on Visual Interfaces to the Social and Semantic Web (VISSW2010)* (2010), pp. 1–5.

[34] COOPER, A. *The inmates are running the asylum*. Sams, 1999.

[35] COSLEY, D., SCHWANDA, V., PEESAPATI, S., SCHULTZ, J., AND BAXTER, J. Experiences with a publicly deployed tool for reminiscing. In *Reminiscence Systems Workshop 2009* (Cambridge, UK, 2009).

[36] CSIKSZENTMIHALYI, M. *Flow: The psychology of optimal experience*. Harper Collins, 1991.

[37] CSIKSZENTMIHALYI, M., AND LEFEVRE, J. Optimal experience in work and leisure. *Journal of Personality and social psychology 56*, 5 (May 1989), 815.

[38] CUNNINGHAM, S., BAINBRIDGE, D., AND FALCONER, A. " More of an art than a science": Supporting the creation of playlists and mixes. In *Proc. ISMIR* (2006).

[39] CUNNINGHAM, S., BAINBRIDGE, D., AND MCKAY, D. Finding new music: a diary study of everyday encounter with novel songs. In *Proceedings of the 8th International Conference on Music Information Retrieval, Vienna, Austria* (2007), pp. 83–88.

[40] CUNNINGHAM, S., DOWNIE, J., AND BAINBRIDGE, D. "The Pain, The Pain": Modelling music information behavior and the songs we hate. In *Proceedings of the 6th International Symposium on Music Information Retrieval* (2005), pp. 474–477.

[41] CUNNINGHAM, S., JONES, M., AND JONES, S. Organizing digital music for use: an examination of personal music. In *ISMIR* (2004).

[42] CUNNINGHAM, S., REEVES, N., AND BRITLAND, M. An ethnographic study of music information seeking: implications for the design of a music digital library. *Proceedings of the 3 rd ACM/IEEE-CS joint conference on Digital Libraries, Houston, USA* (2003).

[43] DACHSELT, R., AND FRISCH, M. Mambo: a facet-based zoomable music browser. In *Proceedings of the 6th international conference on Mobile and ubiquitous multimedia* (2007), ACM, pp. 110—-117.

[44] DENORA, T. *Music in Everyday Life*. Univ. Exeter, Cambridge, UK, 2000.

[45] DIETZ, H.-P. *Supporting association, reminiscing and social discovery in a music player*. Bachelor thesis, Ludwig-Maximilians-Universität, Munich, 2010.

[46] DOWNIE, J. S. Music information retrieval. In *Annual Review of Information Science and Technology 37*, vol. 37. Jan. 2003, pp. 295–340.

[47] DUNN, P. G., DE RUYTER, B., AND BOUWHUIS, D. G. Toward a better understanding of the relation between music preference, listening behavior, and personality. *Psychology of Music* (Mar. 2011), 1–18.

[48] FIELDS, B. *Contextualize Your Listening : The Playlist as Recommendation Engine*. PhD thesis, University of London, UK, 2011.

[49] FIRAN, C. S., NEJDL, W., AND PAIU, R. The Benefit of Using Tag-Based Profiles. In *2007 Latin American Web Conference (LA-WEB 2007)* (Oct. 2007), no. Section 6, IEEE, pp. 32–41.

[50] FLEXER, A., SCHNITZER, D., GASSER, M., AND WIDMER, G. Playlist Generation Using Start and End Songs. In *Proc. ISMIR* (Philadelphia, USA, 2008), pp. 173–178.

[51] FOSTER, A., AND FORD, N. Serendipity and information seeking: an empirical study. *Journal of Documentation 59*, 3 (2003), 321–340.

[52] FREEMAN, E., AND FERTIG, S. Lifestreams: Organizing your electronic life. In *AAAI Fall Symposium: AI Applications in Knowledge Navigation and Retrieval* (1995), pp. 38–44.

[53] GEMMELL, J., BELL, G., AND LUEDER, R. MyLifeBits: a personal database for everything. *Communications of the ACM 49*, 1 (2006), 88–95.

[54] GOLDBERG, D., NICHOLS, D., OKI, B., AND TERRY, D. Using collaborative filtering to weave an information tapestry. *Communications of the ACM 61*, 10 (1992), 1–10.

[55] GOLDER, S. A., AND HUBERMAN, B. A. Usage patterns of collaborative tagging systems. *Journal of Information Science 32*, 2 (Apr. 2006), 198–208.

[56] GOSLING, S. *Snoop: What Your Stuff Says About You*. Basic Books, 2009.

[57] GOUSSEVSKAIA, O., KUHN, M., AND WATTENHOFER, R. Exploring Music Collections on Mobile Devices. In *Proc. 10th international conference on Human computer interaction with mobile devices and services* (2008), ACM, pp. 359–362.

[58] GRAEBSCH, R. *Visualizing Multiple Last.FM Listening Histories*. Bachelor thesis, Ludwig-Maximilians-Universität, Munich, 2010.

[59] GREASLEY, A. E., AND LAMONT, A. Exploring engagement with music in everyday life using experience sampling methodology. *Musicae Scientiae 15*, 1 (Mar. 2011), 45–71.

[60] GREASLEY, A. E., AND LAMONT, A. M. Music preference in adulthood: Why do we like the music we do. In *Proceedings of the 9th International Conference on Music Perception and Cognition (ICMPC9)* (2006), pp. 960–966.

[61] GREWE, O., NAGEL, F., KOPIEZ, R., AND ALTENMÜLLER, E. Listening to Music as a Re-Creative Process: Physiological, Psychological, and Psychoacoustical Correlates of Chills and Strong Emotions. *Music Perception 24*, 3 (2007), 297–314.

[62] HAN, S.-K. Unraveling the Brow: What and How of Choice in Musical Preference. *Sociological Perspectives 46*, 4 (Dec. 2003), 435–459.

[63] HANSEN, D. L., AND GOLBECK, J. Mixing It Up: Recommending Collections of Items. In *Proc. CHI 2009* (2009), ACM, pp. 1217–1226.

[64] HARGREAVES, D., AND NORTH, A. The Functions of Music in Everyday Life: Redefining the Social in Music Psychology. *Psychology of Music 27*, 1 (1999), 71–83.

[65] HARGREAVES, D. J., COMBER, C., AND COLLEY, A. Effects of Age, Gender, and Training on Musical Preferences of British Secondary School Students. *Journal of Research in Music Education 43*, 3 (1995), 242–250.

[66] HARGREAVES, D. J., AND NORTH, A. C., Eds. *The social psychology of music*. Oxford University Press, New York, 1997.

[67] HEER, J., CARD, S., AND LANDAY, J. Prefuse: a toolkit for interactive information visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (2005), ACM, pp. 421–430.

[68] HERRERA, P., RESA, Z. N., AND SORDO, M. Rocking around the clock eight days a week: an exploration of temporal patterns of music listening. In *WOMRAD 2010 - Workshop on Music Recommendation and Discovery* (2010), pp. 7–10.

[69] HITCHNER, S., MURDOCH, J., AND TZANETAKIS, G. Music Browsing Using a Tabletop Display. In *Proc. 8th International Conference on Music Information Retrieval (ISMIR'07)* (Feb. 2007), ISMIR, pp. 175–176.

[70] HODGES, S., WILLIAMS, L., BERRY, E., IZADI, S., SRINIVASAN, J., BUTLER, A., SMYTH, G., KAPUR, N., AND WOOD, K. SenseCam: A retrospective memory aid. In *UbiComp 2006: Ubiquitous Computing* (2006), Springer, pp. 177–193.

[71] HORNBY, N. *High Fidelity*. Penguin Books, London, UK, 2000.

[72] HURON, D. The New Empiricism: Systematic Musicology in a Postmodern Age. Tech. rep., 1999.

[73] HWANG, K.-S., AND CHO, S.-B. Landmark detection from mobile life log using a modular Bayesian network model. *Expert Systems with Applications 36*, 10 (Dec. 2009), 12065–12076.

[74] JENNINGS, D. *Net, Blogs and Rock'n'Roll: How Digital Discovery Works and What It Means for Consumers, Creators and Cultur*. Nicholas Brealey Publishing, 2007.

[75] JOHN, O. P., DONAHUE, E. M., AND KENTLE, R. L. The Big Five Inventory - Versions 4a and 54. Tech. rep., University of California, Berkeley, Institute of Personality and Social Research, Berkeley, CA, USA, 1991.

[76] JONES, S. Popular music, criticism, advertising and the music industry. *Journal of Popular Music Studies 5*, 1 (Mar. 1993), 79–91.

[77] JULIA, C., AND JORDA, S. Songexplorer: a tabletop application for exploring large collections of songs. In *10th International Society for Music Information Retrieval Conference (ISMIR 2009)* (2009), no. Ismir, pp. 675–680.

[78] JUSLIN, P. N., LILJESTRÖM, S., VÄSTFJÄLL, D., BARRADAS, G., AND SILVA, A. An experience sampling study of emotional reactions to music: listener, music, and situation. *Emotion (Washington, D.C.) 8*, 5 (Oct. 2008), 668–683.

[79] KAHNEY, L. *The cult of iPod*. No Starch Press, 2005.

[80] KATZ, E., AND FOULKES, D. On the Use of the Mass Media as "Escape": Clarification of a Concept. *Public Opinion Quarterly 26*, 3 (Jan. 1962), 377–388.

[81] KIM, J.-Y., NATTER, M., AND SPANN, M. Pay what you want: A new participative pricing mechanism. *Journal of Marketing 73*, 1 (Jan. 2009), 44–58.

[82] KINJO, H., AND SNODGRASS, J. G. Does the generation effect occur for pictures? *The American journal of psychology 113*, 1 (Jan. 2000), 95–121.

[83] KNEES, P., SCHEDL, M., POHLE, T., AND WIDMER, G. An innovative three-dimensional user interface for exploring music collections enriched with meta-information from the web. In *Proceedings of the ACM Multimedia* (2006), pp. 17–24.

[84] KNOBLOCH, S., AND ZILLMANN, D. Mood Management via the Digital Jukebox. *Journal of Communication 52*, 2 (June 2002), 351–366.

[85] KOHONEN, T. The self-organizing map. *Proceedings of the IEEE 78*, 9 (1990), 1464–1480.

[86] KONECNI, V. Social Interaction and Musical Preference. In *The psychology of music*, D. Deutsch, Ed. Academic Press, Inc., 1982, pp. 497–516.

[87] KONECNI, V., CROZIER, J., AND DOOB, A. Anger and Expression of Aggression: Effects on Aesthetic Preference. *Scientific Aesthetics 1* (1976), 47–55.

[88] KOSARA, R., HAUSER, H., AND GRESH, D. L. An Interaction View on Information Visualization. In *EUROGRAPHICS* (2003), pp. 123–137.

[89] KREUTZ, G., SCHUBERT, E., AND MITCHELL, L. Cognitive Styles of Music Listening. *Music Perception 26*, 1 (2008), 57–73.

[90] KWONG, T., AND LEE, M. Behavioral intention model for the exchange mode internet music piracy. In *Proceedings of the 35th Hawaii International Conference on System Sciences* (2002), vol. 00, IEEE.

[91] LAMBIOTTE, R., AND AUSLOOS, M. Uncovering collective listening habits and music genres in bipartite networks. *Physical Review E 72*, 6 (2005), 66107.

[92] LAMMING, M., AND FLYNN, M. Forget-me-not: Intimate computing in support of human memory. In *Proceedings of FRIEND21* (1994).

[93] LAMONT, A., AND WEBB, R. Short-and long-term musical preferences: what makes a favourite piece of music? *Psychology of Music 38*, 2 (Aug. 2010), 222–241.

[94] LAPLANTE, A. *Everyday life music information-seeking behaviour of young adults: an exploratory study*. PhD thesis, McGill University, 2008.

[95] LAXMAN, S., AND SASTRY, P. A survey of temporal data mining. *Sadhana 31*, 2 (2006), 173–198.

[96] LAZARSFELD, P. *Radio and the Printed Page*. Duell, Sloan and Pearce, New York, 1940.

[97] LEE, J., AND DOWNIE, J. Survey of Music Information Needs, Uses, and Seeking Behaviours: Preliminary Findings. In *ISMIR Proceedings* (2004), pp. 441–446.

[98] LEHRER, J. *How We Decide*. Houghton Mifflin, New York, 2009.

[99] LEITICH, S., AND TOPF, M. Globe of music-music library visualization using geosom. In *Proceedings of the 8th International Conference on Music Information Retrieval (ISMIR 2007)* (2007), Citeseer, pp. 1–4.

[100] LEONG, T. W., HOWARD, S., AND VETERE, F. Choice: abidcating or exercising? In *Proc. CHI 2008* (2008), pp. 715–724.

[101] LEONG, T. W., VETERE, F., AND HOWARD, S. Randomness as a resource for design. In *Proceedings of the 6th conference on Designing Interactive systems* (New York, New York, USA, 2006), ACM, pp. 132–139.

[102] LEVITIN, D., AND MCGILL, J. Life Soundtracks: The uses of music in everyday life. Tech. Rep. 514, 2007.

[103] LEVITIN, D., AND TIROVOLAS, A. Current advances in the cognitive neuroscience of music. *Annals of the New York Academy of Sciences 1156*, 1 (Mar. 2009), 211–231.

[104] LEWIS, J. IBM computer usability satisfaction questionnaires: psychometric evaluation and instructions for use. *International journal of human computer interaction 7*, 1 (1995), 57–78.

[105] LI, I., DEY, A., AND FORLIZZI, J. A stage-based model of personal informatics systems. In *Proceedings of the 28th international conference on Human factors in computing systems* (New York, New York, USA, 2010), ACM, pp. 557–566.

[106] LILLIE, A. S. *MusicBox: Navigating the space of your music*. Master's thesis, MIT, Boston, MA, USA, 2009.

[107] MANN, S. Smart Clothing : The Wearable Computer and WearCam. *Personal Technologies 1*, 1 (Mar. 1997), 21–27.

[108] MENON, V., AND LEVITIN, D. J. The rewards of music listening: response and physiological connectivity of the mesolimbic system. *NeuroImage 28*, 1 (Oct. 2005), 175–184.

[109] MERRIAM, A. P. *The Anthropology of Music*, 1st ed. Northwestern University Press, 1964.

[110] MILLIMAN, R. Using Background Music To Affect Behavior of Supermarket Shoppers. *The Journal of Marketing 46*, 3 (1982), 86–91.

[111] MOD, C. *Post-Artifact Books and Publishing*. Online, 2011.

[112] MUNZNER, T. A nested model for visualization design and validation. *IEEE transactions on visualization and computer graphics 15*, 6 (2009), 921–928.

[113] NEUMAYER, R., DITTENBACH, M., AND RAUBER, A. Playsom and pocketsomplayer, alternative interfaces to large music collections. In *Proc. ISMIR* (2005), pp. 618–623.

[114] NIELSEN. The hyper-fragmented world of music: Marketing considerations and revenue maximisation. Tech. Rep. March, Nielsen Music, New York, 2011.

[115] NORTH, A., AND HARGREAVES, D. The effects of music on responses to a dining area. *Journal of Environmental Psychology 16*, 1 (1996), 55–64.

[116] NORTH, A. C., AND HARGREAVES, D. J. Musical preferences during and after relaxation and exercise. *The American journal of psychology 113*, 1 (2000), 43–67.

[117] NORTH, A. C., AND HARGREAVES, D. J. Lifestyle correlates of musical preference: 1. Relationships, living arrangements, beliefs, and crime. *Psychology of Music 35*, 1 (Jan. 2007), 58–87.

[118] NORTH, A. C., AND HARGREAVES, D. J. Lifestyle correlates of musical preference: 2. Media, leisure time and music. *Psychology of Music 35*, 2 (Feb. 2007), 179–200.

[119] NORTH, A. C., AND HARGREAVES, D. J. Lifestyle correlates of musical preference: 3. Travel, money, education, employment and health. *Psychology of Music 35*, 3 (July 2007), 473–497.

[120] NORTH, A. C., HARGREAVES, D. J., AND HARGREAVES, J. J. Uses of Music in Everyday Life. *Music Perception 22*, 1 (Sept. 2004), 41–77.

[121] NORTH, A. C., HARGREAVES, D. J., AND MCKENDRICK, J. The Influence of In-Store Music on Wine Selections. *Journal of Applied Psychology 84*, 2 (1999), 271–276.

[122] OBERHOLZER-GEE, F., AND STRUMPF, K. The Effect of File Sharing on Record Sales: An Empirical Analysis. *Journal of Political Economy 115*, 1 (Feb. 2007), 1–42.

[123] OLIVER, N., AND FLORES-MANGAS, F. MPTrain: a mobile, music and physiology-based personal trainer. In *Proc. MobileHCI '06* (2006), ACM, pp. 21–28.

[124] PAMPALK, E. Islands of music: Analysis, organization, and visualization of music archives. *Journal of the Austrian Soc. for Artificial Intelligence 22*, 4 (2003), 20–23.

[125] PAMPALK, E., POHLE, T., AND WIDMER, G. Dynamic Playlist Generation Based on Skipping Behaviour. In *Proc. ISMIR* (2005), pp. 634–637.

[126] PAMPALK, E., RAUBER, A., AND MERKL, D. Content-based organization and visualization of music archives. In *Proceedings of the tenth ACM international conference on Multimedia - MULTIMEDIA '02* (New York, New York, USA, 2002), ACM Press, pp. 570–579.

[127] PARK, C. H., AND KAHNG, M. Temporal Dynamics in Music Listening Behavior - A Case Study of Online Music Service. In *Proc. IEEE/ACIS* (Aug. 2010), IEEE, pp. 573–578.

[128] PAUWS, S., AND WIJDEVEN, S. V. D. User evaluation of a new interactive playlist generation concept. In *Proc. ISMIR* (2005), pp. 638–643.

[129] PEITZ, M., AND WAELBROECK, P. The effect of internet piracy on music sales: cross-section evidence. *Review of Economic Research on Copyright Issues 1*, 2 (2004), 71–79.

[130] PERLIN, K., AND FOX, D. Pad - An Alternative Approach to the Computer Interface. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques* (1993), ACM, pp. 57–64.

[131] PLAISANT, C. The challenge of information visualization evaluation. In *Proceedings of the working conference on Advanced visual interfaces* (2004), vol. 11, ACM, pp. 109–116.

[132] PLAISANT, C., MILASH, B., ROSE, A., WIDOFF, S., AND SHNEIDERMAN, B. LifeLines: visualizing personal histories. In *Proceedings of the SIGCHI conference on Human factors in computing systems common ground - CHI '96* (New York, New York, USA, 1996), ACM Press, pp. 221–ff.

[133] POUSMAN, Z., STASKO, J., AND MATEAS, M. Casual Information Visualization: Depictions of Data in Everyday Life. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (2007), 1145–1152.

[134] RAGNO, R., BURGES, C. J. C., AND HERLEY, C. Inferring similarity between music objects with application to playlist generation. In *Proc. 7th ACM SIGMM international workshop on Multimedia information retrieval* (2005), ACM Press New York, NY, USA, pp. 73–80.

[135] RENTFROW, P. J., GOLDBERG, L. R., AND LEVITIN, D. J. The structure of musical preferences: A five-factor model. *Journal of personality and social psychology 100*, 6 (June 2011), 1139–57.

[136] RENTFROW, P. J., AND GOSLING, S. D. The do re mi's of everyday life: The structure and personality correlates of music preferences. *Journal of Personality and Social Psychology 84*, 6 (2003), 1236–1256.

[137] RENTFROW, P. J., AND GOSLING, S. D. Message in a ballad: the role of music preferences in interpersonal perception. *Psychological science 17*, 3 (Mar. 2006), 236–242.

[138] RENTFROW, P. J., AND GOSLING, S. D. The content and validity of music-genre stereotypes among college students. *Psychology of Music 35*, 2 (Feb. 2007), 306–326.

[139] ROB, R., AND WALDFOGEL, J. Piracy on the high C's: Music downloading, sales displacement, and social welfare in a sample of college students. *The Journal of Law and Economics 49*, 1 (Apr. 2004), 29–62.

[140] ROE, K. SWEDISH YOUTH AND MUSIC: Listening Patterns and Motivations. *Communication Research 12*, 3 (July 1985), 353–362.

[141] ROEDIGER, H., AND MCDERMOTT, K. Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition 21*, 4 (1995), 803–814.

[142] RÖSING, H. Listening behaviour and musical preference in the age of 'transmitted music'. *Popular Music 4* (Nov. 1984), 119–149.

[143] ROSS, P. Grundlagen einer musikalischen Rezeptionsforschung. In *Rezeptionsforschung in der Musikwissenschaft*, H. Rösing, Ed. Wissenschaftliche Buchgesellschaft, Darmstadt, Germany, 1983, pp. 377–418.

[144] SCHWARTZ, B. *The paradox of choice: Why more is less*. Harper Perennial, New York, 2005.

[145] SCHWARTZ, K. D., AND FOUTS, G. T. Music preferences, personality style, and developmental issues of adolescents. *Journal of Youth and Adolescence 32*, 3 (2003), 205–213.

[146] SEASE, R., AND MCDONALD, D. W. Musical fingerprints: Collaboration around home media collections. In *Proceedings of the ACM 2009 international conference on Supporting group work* (2009), ACM, pp. 331–340.

[147] SEIFFERT, F. *Visualizing Music Listening Histories*. PhD thesis, Ludwig-Maximilians-Universität, München, 2010.

[148] SELLEN, A., FOGG, A., AITKEN, M., HODGES, S., ROTHER, C., AND WOOD, K. Do life-logging technologies support memory for the past?: an experimental study using sensecam. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (2007), ACM, pp. 81–90.

[149] SELLEN, A., AND WHITTAKER, S. Beyond total capture: a constructive critique of lifelogging. *Communications of the ACM 53*, 5 (2010), 70–77.

[150] SHIRKY, C. Ontology is Overrated: Categories, Links, and Tags. Tech. rep., 2005.

[151] SHNEIDERMAN, B. The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages* (1996), IEEE Comput. Soc. Press, pp. 336–343.

[152] SIEK, K. A., ROGERS, Y., AND CONNELLY, K. H. Fat finger worries: how older and younger users physically interact with PDAs. In *INTERACT 2005* (2005), pp. 267–280.

[153] SILFVERBERG, S., LIIKKANEN, L., AND LAMPINEN, A. I'll press play, but I won't listen: profile work in a music-focused social network service. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work* (2011), ACM, pp. 207–216.

[154] SLOBODA, J. A., O'NEILL, S. A., AND IVALDI, A. Functions of music in everyday life: An exploratory study using the Experience Sampling Method. *Musicae Scientiae 5*, 1 (2001), 9–32.

[155] STAVNESS, I., GLUCK, J., VILHAN, L., AND FELS, S. The MUSICtable: A Map-Based Ubiquitous System for Social Interaction with a Digital Music Collection. *Lecture Notes in Computer Science 3711/2005*, Entertainment Computing - ICEC 2005 (June 2005), 291–302.

[156] STOBER, S., AND NÜRNBERGER, A. Towards User-Adaptive Structuring and Organization of Music Collections. In *Adaptive Multimedia Retrieval. Identifying, Summarizing, and Recommending Image and Music*. Springer, Berlin, Germany, 2008, pp. 53–65.

[157] STOBER, S., AND NÜRNBERGER, A. A multi-focus zoomable interface for multi-facet exploration of music collections. *7th International Symposium on Computer Music Modeling and Retrieval* (2010), 339–354.

[158] SUROWIECKI, J. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Anchor, New York, NY, USA, 2004.

[159] THOMPSON, R. Social context and the subjective experience of different types of rock music. *Journal of youth and adolescence 24*, 6 (Dec. 1995), 731–744.

[160] TORRENS, M., HERTZOG, P., AND ARCOS, J. Visualizing and exploring personal music libraries. In *Proc. 5th International Conference on Music Information Retrieval (ISMIR '04)* (2004), ISMIR, pp. 421–424.

[161] TUFTE, E. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, USA, 1983.

[162] TULVING, E. Episodic and semantic memory. In *Organization of Memory*. 1972, pp. 382–402.

[163] van De Wetering, H., van Gulik, R., and Vignoli, F. Mapping music in the palm of your hand, explore and discover your collection. In *Proc. ISMIR* (2004), pp. 409–414.

[164] van Gulik, R., and Vignoli, F. Visual playlist generation on the artist map. In *Proc. ISMIR* (2005), vol. 2005, pp. 520–523.

[165] Viégas, F., Boyd, D., Nguyen, D. H., Potter, J., and Donath, J. Digital Artifacts for Remembering and Storytelling: PostHistory and Social Network Fragments. In *Proceedings of the 37th Hawaii International Conference on System Sciences* (2004), vol. 00, pp. 1–10.

[166] Viégas, F., Golder, S., and Donath, J. Visualizing email content: portraying relationships from conversational histories. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (2006), ACM, pp. 979–988.

[167] Vignoli, F. Digital Music Interaction concepts: a user study. In *Proc. ISMIR* (2004).

[168] Vogel, D., and Baudisch, P. Shift: a technique for operating pen-based interfaces using touch. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, 2007), no. c, ACM, pp. 657–666.

[169] Voida, A., Grinter, R. E., Ducheneaut, N., Edwards, W. K., and Newman, M. W. Listening in: practices surrounding iTunes music sharing. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (2005), ACM, pp. 191–200.

[170] Voida, S., Tobiasz, M., Stromer, J., Isenberg, P., and Carpendale, S. Getting Practical with Interactive Tabletop Displays: Designing for Dense Data, "Fat Fingers," Diverse Interactions, and Face-to-Face Collaboration. In *Proc. ACM International Conference on Interactive Tabletops and Surfaces (ITS)* (2009), ACM, pp. 109–116.

[171] Ward, D., Blackwell, A., and MacKay, D. Dasher—a data entry interface using continuous gestures and language models. In *Proceedings of the 13th annual ACM symposium on User interface software and technology* (2000), ACM New York, NY, USA, pp. 129–137.

[172] Ward, D. J., and MacKay, D. Fast Hands-free Writing by Gaze Direction. *Nature 838*, August (2002), 4–6.

[173] Ware, C. *Information Visualization: Perception for Design*, 2nd editio ed. Morgan Kaufmann, 2004.

[174] Ware, C., Purchase, H., Colpoys, L., and McGill, M. Cognitive measurements of graph aesthetics. *Information Visualization 1*, 2 (June 2002), 103–110.

[175] WATTENBERG, M. Arc diagrams: visualizing structure in strings. In *IEEE Symposium on Information Visualization, 2002. INFOVIS 2002.* (2002), IEEE Comput. Soc, pp. 110–116.

[176] WILLS, S. A., AND MACKAY, D. J. C. DASHER–an efficient writing system for brain-computer interfaces? *IEEE transactions on neural systems and rehabilitation engineering 14*, 2 (June 2006), 244–246.

[177] YALCH, R., AND SPANGENBERG, E. Effects of store music on shopping behavior. *Journal of Consumer Marketing 7*, 2 (1993), 55–63.

[178] ZHELEVA, E., GUIVER, J., MENDES RODRIGUES, E., AND MILIĆ-FRAYLING, N. Statistical models of music-listening sessions in social media. In *Proc. WWW '10* (New York, New York, USA, 2010), ACM Press, pp. 1019–1028.

[179] ZHOU, M., AND FEINER, S. Visual task characterization for automated visual discourse synthesis. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (New York, New York, USA, 1998), ACM Press/Addison-Wesley Publishing Co., pp. 392–399.

[180] ZILLMANN, D. Mood Management: Using Entertainment to Full Advantage. In *Communication, social cognition, and affect*, L. Donohew, H. E. Sypher, and E. T. Higgins, Eds. 1988, pp. 147–171.

[181] ZWEIGENHAFT, R. L. A Do Re Mi Encore. *Journal of Individual Differences 29*, 1 (Jan. 2008), 45–55.